A COMPUTATIONAL SEMANTICS FOR ASPECT — DRAFT

by

Srini Narayanan
University of California at Berkeley
and International Computer Science Institute

September 28, 1994

Abstract

This report describes a new computational model for modeling the semantics of aspect. The model described suggests an ontology for situations based on a *natural* logic of processes. Aspect is modeled by an *active* model, where the semantics of aspect, including possible entailments and inferences can be obtained by inspecting the executing model. This allows the model to provide compact solutions to some well known problems in previous attempts. The report also attempts to demonstrate the sufficiency of the model to model aspectual distinctions made by the world's languages. Finally, the proposed approach seems to have some promise in modeling narrative interpretation. The rest of the report describes the model in detail to justify these claims.

Contents

1 Introduction

Aspects allow a single situation to be viewed in multiple ways by focusing on the *internal temporal character* of the situation. Situations are are composed of constituent sets of phases (such as beginnings, middles, ends, and whether the situation is the result of several iterations of an event). Aspectual words, such as *start*, *end*, *cease*, *continue*, and *stop* exist in all languages and allow different perspectives on a single situation. [1]

Aspect Markers refer to Inflexional affixes, Auxiliaries, and other **Grammatical Constructions** that *select* components of situations. Aspect Markers have been identified and studied in many languages [Comrie1976]. In English, for instance, the *Progressive* aspect is marked by the *-ing* suffix as in *playing*, and the *Habitual* by the Construction *Used To* (as in *He used to play baseball on Sundays*). Russian, and other Slavonic languages make extensive use of Inflectional Affixes to mark aspect.

While Aspect Markers are found across languages, what makes Aspect a particularly interesting phenomena is that in all languages studied, the *natural or inherent* semantics (commonsense knowledge) of the situations described combine with and modify the interpretations and entailments of the Aspect Marker System. This rather close coupling between Commonsense Knowledge and Linguistic Constructions has been recognized by most researchers of aspect. This has also made aspect a subject of philosophical and linguistic study since Aristotle.

Linguists and Philosophers of Language who have attempted to give an account of aspect have distinguished the use of Aspect into two categories, namely the

1. **Aspectual Class** of a predicate.
2. **Aspectual Form** of a predicate.

Aspectual class refers to the Aristotelian Class [Vendler1967] to which a basic verb belongs, and Aspectual Form refers to the specific Marker System the verb occurs with. The seminal work on modeling the semantics of Aspect was done by Dowty [Dowty1979]. In his account, Dowty uses Model Theoretic Semantics (MTS) to look at truth-functional relations between semantic representations of Aspect. Using time intervals as an ontological primitive, Dowty's MTS models aspect by looking at truth-values *entailed* by different situational phases (relative to time intervals). One major problem with this kind of analysis is that aspectual distinctions do not lead to differences in truth-values unless the speaker's perspective is modeled as well. This problem gives rise to several paradoxes and complications in the Model Theoretic Account.

More recently, Computational Linguists have become interested in modeling Aspect. The need for encoding descriptions of complex events and dynamic processes, as well as making inferences and computing entailments in a tractable manner makes it a natural problem for Computational Linguists. An influential attempt at modeling aspect can be found in the work of Moens and

[1] Aspects differ from Tenses in that whereas tenses deal with the temporal relations between situations (such as past, present and future), aspects enable focus on the compositional attributes of a situation. So while, tense and aspect interact with each other to generate the overall meaning of an utterance, they are usually considered distinct entities by most researchers.

Steedman [Moens1988] who convincingly argue that aspect requires an ontology that is based on the logic of processes rather than temporal primitives such as time intervals or time points. They propose a basic tripartite representational structure called the **nucleus**[2] which can be thought of as a core ontology of processes. They also propose a *Type Coercion* network based on the **nucleus** which controls how *Aspectual Class* interacts with *Aspectual Form*. Others [Dorr1993] [Webber1988] have demonstrated the usefulness of this idea in modeling narrative interpretation and in machine translation. One problem in the Moens and Steedman's model of aspect is the requirement of a special set of *type coercion* meta-rules whose generality and motivatedness is unclear.

The work described here proceeds from the assumption (as with Moens and Steedman) that any model of Aspect requires an ontology of processes and events. However, the model described here differs from all previous attempts at modeling aspect in satisfying the following two basic principles of cognitive modeling.

1. **Strong Operationality** principle.
2. **Naturalness** principle.

These principles are elaborated below.

PRINCIPLE: Strong Operationality Principle

Operationality in a strong sense means that the semantic representation of aspect should grounded in the execution of the model itself, thereby precluding the need for any special model interpretation rules.

PRINCIPLE: Naturalness Principle

The **Naturalness Principle** corresponds to satisfying the following conditions.

1. The primitive elements of the proposed ontology should be natural distinctions of our experience of processes and events in the world.
2. Relevant inferences and aspectual entailments should fall out of the executing model without special *type coercion* or other meta-rules.
3. The model proposed here should deal with cross-linguistic data, suggesting the possibility of a natural ontology for aspect.

Dowty's Model Theoretic Semantics fails both principles. Most of the more recent computational models (including the Moens and Steedman model) satisfy the first condition of the **Naturalness** principle. No computational model of aspect satisfies the **Strong operationality** principle.

[2] The nucleus consists of a causal sequence of **preparatory process - culmination point - consequences** which provides a total ordering of situational stages.

In the model described here, aspect is characterized using a set of natural ontological primitives that are composed into an **active** model. It is the active property of the proposed model that allows the model to satisfy the **Strong Operationality** principle. This property also allows for an elegant solution to the **Imperfective Paradox** [Dowty1979] and other problems. I also claim that the proposed ontology satisfies the conditions for the **Naturalness** principle.

1.1 Summary

In summary, the aspect is an interesting problem for Knowledge Representation and Inference because of the following reasons.

1. The problem of Aspect requires the modeling of complex events and situations in a dynamic world.
2. Aspect requires the integration of commonsense knowledge (Lexical Semantics) with Linguistic knowledge (Constructions).
3. Any computational semantics of Aspect has to be *operationalizable*, ie. have some means of modeling the processing level issues involved.
4. Modeling Aspect is a prerequisite to modeling human interpretations of narratives.

The model of aspect described in this report makes the following novel contributions.

1. We propose a new set of parameters to model Aspectual Form of a language as well the Aspectual Class of situations. These parameters suggest a natural ontology for aspect based on the logic of processes.
2. In our model, the meaning of a linguistic input results from the *continuous*, *on-line* interaction between input and memory. In the case of aspect, we demonstrate how such an active interpretation model solves many problems and paradoxes in earlier accounts.
3. The model of Aspect proposed explains cross-linguistic data.

The rest of this report describes the model in detail in an attempt to demonstrate these claims. First the linguistic theory of aspect is discussed pointing out the various distinctions made by the Aspect Marker system of the world's languages. The report then discusses an ontology of situations including a new *feature semantics* for the Aspectual Class of a predicate. Once the ontological primitives have been discussed, the *active* computational model of aspect is described along with examples demonstrating its sufficiency for cross-linguistic modeling. The report then turns to modeling narrative inferences. We close with some discussion of a possible connectionist implementation of our model, as well as some future directions for research.

2 Linguistic Theory

In this section, we briefly describe the current linguistic theory regarding aspect. Individual parts of the theory will be dealt with in greater detail in subsequent sections. The basic distinctions that characterize aspect are divided into the following three types.

1. The Aspect Marker System.
2. Aspectual Verbs, Adverbs, and Extended Aspect.
3. Lexical Semantics.

In this section we will deal with the first two categories. The next two sections deal with Lexical semantics.

2.1 The Aspect Marker System

Figure 1 shows a language independent characterization of the Aspect Marker system. The data is taken from [Comrie1976]. Aspect Marker Systems (different languages) mark different subsets of the tree shown. The basic distinctions are described below.

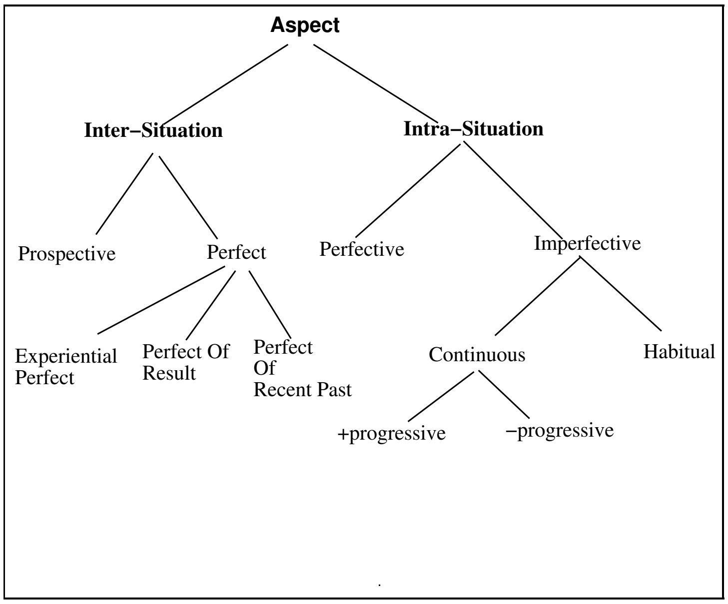

Figure 1: The Aspectual Taxonomy.

- **Intra-Situation Vs. Inter-Situation** An Aspect Marker can either refer to the internal temporal consistency of a single situation by focusing on different logical stages of a situation or relate some state to a preceding situation. Consider the following two examples in English.

Example 1: John was playing tennis.

Example 2: John has been playing tennis.

While Example 1 refers to the process of playing tennis (one part of the play tennis situation), Example 2 refers to the *continuing relevance* of the play tennis situation. For instance, in Example 2, a reasonable inference might be *John is tired now.* While there is a debate about whether inter-situational relevance should be considered aspect or tense, this report will deal with both categories as aspectual categories.

- **Inter-situational** In the inter-situational case, **Perfect** aspect refers to the continuing relevance of a past situation, whereas **Prospective** refers to the future relevance of some state. Example 2 is an example of the Perfect aspect. In English, the Prospective aspect is marked by constructions such as *about to*, *to be going to*, and *on the point of.*
- **Intra-Situational**

  Intra-situational aspect consists of the following two categories, namely

  1. Perfective Aspect.
  2. Imperfective Aspect.

  These two categories are discussed below

  * **Perfective Aspect**

    Perfective aspect refers to a situation as a complete whole. Perfective aspect is defined by its *lack* of reference to the internal temporal consistency of a situation. Russian has special affixes (such as *pos*) to denote the Perfective aspect. English does not have a special construction to mark the Perfective, and the Perfective/Simple Past distinction in English is context dependent.
  * **Imperfective Aspect**

    Imperfective aspect refers to a the internal state of a situation; namely viewing a situation from within. In Example 1, we see the described situation as one in progress in the past. Such a perspective is referred to as the *Progressive* aspect. Consider the contrast with the sentence

    Example 3: John used to play tennis.

    Here the *used to* construction suggests the *habitual* aspect in interpreting the Example 3. The habitual aspect refers to the iteration of a process multiple times. Habitual aspect can combine with the progressive as in *John used to be playing tennis.* We will examine these distinctions in greater detail in Section 7.

2.2 Other ways Of Expressing Aspect

While most languages have grammatically obligated structures for expressing aspect, there may also be many optional devices used. These techniques generally fall into four main categories [Berman1994] These methods include:

1. *Morphological marking* on the verb such as particles in English (as in eat up, fall down).

2. *Verbs of Aspect*, such as start, finish, stop, etc.

3. *Adverbial phrases* such as on the verge of, almost, again and again.

4. *Iteration* as in run and run.

The table below illustrates some of the ways of expressing aspects in two different languages, namely in *Hindi* and English. These categories have been shown to be present in most world languages [Berman1994].

Aspectual Expression	**Hindi Example**	**English Example**
Achievement	*sufal*	succeed
Cessive	*rok*	stop
Completive	*khatm*	finish
Cumulative	*ek ke baad ek*	one by one
Imminent	None	Almost, nearly
Inceptive/Inchoative	*shuru*	begin, start
Iterative	*chal chal*	*walk and walk*
Perfect	*abhi tak* (until now)	still
Protracted	*suffix -te* as in *chalte*	keep on

Given the basic linguistic theory of aspect, we now turn to an ontology of situations that can adequately capture some of the distinctions mentioned above.

3 Ontological Primitives

This section outlines the ontology of situations that model aspect. As was discussed in the introduction, the semantics of aspect results from the interaction of a linguistic Aspect Marker system with the natural semantics of the described situation. Our ontology reflects this distinction and the discussion proceeds in two parts.

1. **Composition Of Situations**.
2. **A Feature Semantics for Aspectual Class**.

Let us consider each of these items in detail.

3.1 Composition Of Situations

This work assumes that situations are composed of the component *causal graph* shown in Figure 2 (a).

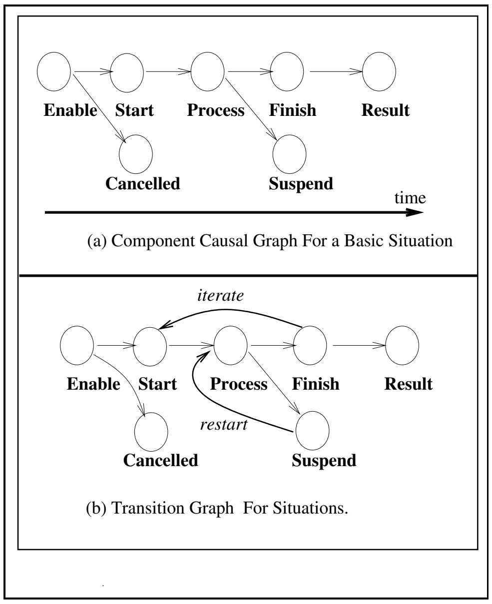

Figure 2: Stages Of Situations, (a) shows the basic component graph, and(b) shows the various possible transitions.

Most elements in this ontology (except the **Suspend** and **Cancelled**) stages have been proposed in models of processes by others [Nakhimovsky1988] and are fairly uncontroversial. The *Suspend* stage models the cessation of processes due to some *stop* interruption. The *Cancelled* stage models the case where a situation is enabled but never obtained.

All stages in the *causal graph* are *also* independently motivated from Linguistics in that each of these stages can be directly referred to by words in the lexicon. The three possible total orderings of the *causal graph* create *three* linear temporal orders for the evolution of situations. These total orders are

1. *Enable-Cancelled.*
2. *Enable-Start-Process-Finish-Result.*
3. *Enable-Start-Process-Suspend.*

Figure 2(b) shows the finite state component *transition graph* that models the possible transitions in modeling situations. The presence of loops makes this graph non-causal. In the *transition graph*, the additional arcs *iterate* and *restart* have obvious interpretations. A situation evolves in stages that follows the transition graph. An actual trace through the *transition graph* is referred to as an **episode**.

Here again, independent linguistic evidence exists for these paths in the graph.

Grammatical Aspect *focuses* on specific stages or sets of stages of the *Transition Graph.* For instance, the *Imperfective* aspect focuses on the *Process* stage of the transition graph, while *Habitual* Aspect Markers pick the *Start-Process-Finish-Iterate* loop.

Our computational model will be an operational model based on the simple transition graph shown in Figure 2. Any direct implementation of the transition graph will capture many effects of the Aspect Marker System, but will fail to address the ways in which the natural semantics of of a situation interacts with the Aspect Marker System. In order to characterize this interaction, we now turn to the issue of the basic ontology of situations.

3.2 A Feature Semantics For The Aspectual Class Of Situations

The Aspectual Class of a Predicate is determined by the following **features**.

1. Duration.
2. Periodicity.
3. Effort.
4. Resources.
5. Agency.

These features are discussed in detail below.

- **Duration** Most situations have durations that fall into specific **time scales**. The table below shows some typical situations along with the associated time scales.

Situation	**Duration**
walk	minutes
run	minutes
cough	second
know	years
read a book	hours
read *War and Peace*	days
sleep	hours

Note that the precision with which we reason about the durations of activities seems to cut the continuous time line into discrete time scales. Note also that the default time scale can be overridden by specific information. For instance, the default duration of reading a book can be overridden if we know the book to be of a specific size.

We will have more to say about the values and uses of the time scale idea in aspect processing in the next section.

- **Periodicity**

  Talmy [Talmy1985] characterizes Verb Aspect as the pattern of distribution of action through time. His characterization classifies verbs based on the distributivity properties of the semantics of verbs. The following classification scheme is a generalization of the scheme of Talmy. The set of verbs can be classified in the following categories.

  1. Certain predicates refer to processes that are inherently *Aperiodic*. Such processes may be due to some irreversible terminal condition obtaining (as in the case of the verb *die*), or may be due to some world state that has to be reversed before the process can be initiated again (as in the case of *fall*).
  2. Certain predicates refer to processes that are inherently *Periodic*. At the end of each period, the conditions that obtain in the world are such that the process is re-enabled. The default periodicity of the processes may be 1 (as in the case of *flash* or *hit*) or > 1 (as in *breathe*).
  3. Certain predicates refer to processes that are composed of specific *sequence* of a other activities. Consider, for instance the predicate *waltz*. This minimally requires a sequence of three patterns of motion (specific length and direction of steps). The overall activity is periodic in that by default it is iterated multiple times.

- **Effort**

  Effort is the expending of energy to maintain an activity.

  A well-motivated distinction that can be made between different situations is between those that require the use of *energy* and those that don't. Our experience of activities such as *running*, *walking*, etc. makes the distinction an extremely salient one, since if we don't expend energy in these activities they terminate. In other cases, such as *knowing something* (or *believing something*), the situation continues despite the absence of continuous effort.

- **Resource**

Generalizing from the idea of effort, one could obtain the more general notion of a *resource requirement* in situation. Consider the situation referred to by the sentence *John is reading a book.* Apart from the effort required (attention, scanning, etc.) there is also an inherent resource consumption where the unread book is the resource and the reading is the consumption of the resource. So the activity described obtains till the book is read (the total time being determined by parameters such as size, and reading rate/per unit size). Note that the resource here is renewable (in that the book can be re-read).

Resources can be classified into the following three categories:

1. World Conditions that hold in a situation. This includes the classic STRIPS like *preconditions* as well as the locking and subsequent release of shared resources.
2. World objects that are *consumed* in the situation. While **energy** is a special type of resource that falls into this category, there may be other consumable resources that characterize a situation (for instance *food* in an *eating* situation).
3. World objects that are *created* or *destroyed* in the situation. Often the situation is described as **telic** when a special object is created at the end of the situation. Situations like *make a chair* and *demolish the building* require the explicit creation or destruction of an object.

- **Agency**

  Agency is a commonly used parameter, which can modify the Aspectual Class of a situation. A situation can be *intentional* or *unintentional.* All intentional situations are *controllable*, whereas unintentional situations can be *controllable* or *uncontrollable.* [3]

Before proceeding further, we need one abstract parameter that is derivable from the parameters mentioned earlier and the context. This parameter pertains to the relationship of the time scale of an activity to the time scale of reasoning.

3.3 Reasoning in Multiple Time Scales

In this section, we will formally define a method of abstraction which models contextual interaction in defining Verb Classes. The principle is quite simple and is stated below:

Definition : The Multiple Time Scale Hypothesis

The interaction between a temporal granularity of interest (called the Reference Time Scale) and the temporal granularity of a described situation (called the Situation Time Scale) is used to create a *Temporal Abstraction Hierarchy*, where levels in the hierarchy control the precision with which the given situation is represented.

Several attempts at modeling commonsense knowledge have led AI researchers to propose that human perception/reasoning of activities deals with specific time scales (such as *Lifetimes* [Mcdermott1982]

[3] Intentionality and controllablilty interact with the *force dynamic* characteristics of situations. The set of ontological primitives for force dynamics is outside the scope of this report.

or Grain [Hobbs1985]). Nakhimovsky [Nakhimovsky1988] points out that such concepts have a natural way of cutting the continuous time parameter space into discrete units. He proposes that the temporal relations **much greater than** and **greater-than** be used in combination with natural units (seconds, minutes hours, days, weeks, months, years and multiples of 5 for each of these units) to control the granularity of reasoning. For instance, he points out the difference in the precision of reasoning in the two cases shown below

Example 4: I'll be back in five minutes.

Example 5: I'll be back in six minutes.

In the first case, a much larger deviation from 5 minutes (a minute) can be tolerated than in the second (a minute cannot be tolerated).

Time scales set up a way to partition verbs into aspectual classes. Using natural time scales and the the relation **much greater than** ($\gg$), I propose an automatic method to construct a temporal abstraction hierarchy to control the precision of reasoning.

The temporal abstraction hierarchy deals with both **periodic** and **aperiodic** processes. Agents may reason at different levels of the temporal hierarchy and consequently arrive at different conclusions about the same situation. Hence, one would expect the choice of time scale to be *subjective*, and so the speaker always has to specify the additional disambiguation information.

However, our shared experience of the world allows the speaker to assume that the hearer will associate default values for the time scale of reasoning based on **natural** units of the hierarchy. Given such natural units, we use the following relevant parameters to construct an *objective* reasoning strategy R.

1. A reference level of reasoning characterized by the natural time scale s_R.
2. A given situation/process which operates on a given time scale s_i. Let $p_{s_i}(\theta)$ be the relevant situation, where θ summarizes the state variables.
3. The temporal relation $\gg$ (**much greater than**).

The strategy is *objective* in the sense that the same reasoning strategy is used by different agents to control the precision of reasoning for a *given* reference time scale s_R.

Equations (1) - (5) below set up this hierarchy formally for different types of situations.

- **Aperiodic Processes**

  The basic idea is if we are interested at a certain time scale s_i, we consider as **Processes**, activities that happen on that time scale ($p_{s_i}(\theta)$). Activities on much smaller time scales are assumed to be instantaneous and qualify as **Events**. Activities on much larger scales are assumed to be unchanging over the episode of reasoning and can be referred to as **Unchanging States**.

  Let $p(\theta)$ be an Aperiodic process whose time scale (start to finish) is s_i. and s_R is a reference time scale of reasoning.

The following equations set up the natural hierarchy for aperiodic processes.

$$R : \forall p_{s_i}(\theta) : s_R \gg s_i \Rightarrow \frac{\delta p(\theta)}{\delta t} = \infty. \tag{1}$$

$$R : \forall p_{s_i}(\theta) : (s_R \not\gg s_i) \wedge (s_i \not\gg s_R) \Rightarrow -\infty < \frac{\delta p(\theta)}{\delta t} < \infty. \tag{2}$$

$$R : \forall p_{s_i}(\theta) : s_i \gg s_R \Rightarrow \frac{\delta p(\theta)}{\delta t} = 0. \tag{3}$$

- **Periodic Processes**

  In the case of periodic process, our clustering will have to be done on two axes, namely the frequency and the duration axes. If the Reference time scale of interest is s_R, and the time scale of the process $p_{s_i}(\theta)$ is s_i, and the periodicity parameter is f ($f \in [single, multiple]$); we have the following possible cases:

  1. Short Duration, Multiple Process.
  2. Short Duration, Single Process.
  3. Long Duration Process.

  Long duration processes satisfy $s_i \gg s_R$, and **Short** duration processes satisfy $s_R \gg s_i$. Processes that are not **Short** or **Long** duration can be considered aperiodic for the purposes of this discussion, and satisfy equation (2). The interesting case for a periodic process is when the described situation is of a short duration, but can be repeated multiple times. Equation (4) shows how this case gets transformed into a continuous process resulting from the *integration* of the multiple occurrences. Equation (5) shows that if the periodic process is of a very long duration, the process gets transformed to a *state*.

  1. *Short Duration, Multiple* Processes.

  $$R : \forall p_{s_i}(\theta) : (s_R \gg s_i) \Rightarrow -\infty < \frac{\delta \int p(\theta)\delta\theta}{\delta t} < \infty. \tag{4}$$

  2. *Short Duration, single* Processes.

  $$R : \forall p_{s_i}(\theta) : (s_R \gg s_i) \Rightarrow \frac{\delta p(\theta)}{\delta t} = \infty. \tag{5}$$

  3. *Long Duration* Processes.

  $$R : \forall p_{s_i}(\theta) : (s_i \gg s_R) \Rightarrow \frac{\delta p(\theta)}{\delta t} = 0. \tag{6}$$

To discuss how the scales of reasoning affect the interpretation, consider the following examples

1. This morning, I was playing tennis.
2. In 1993, I was playing tennis.

Clearly in (1), most people would interpret the situation as referring to a single instance of playing tennis, whereas in (2) the interpretation would be one which refers to several instances of playing tennis (similar to the habitual interpretation *I used to play tennis*). Why are the two interpretations different?

The only difference between (1) and (2) is in the reference time scale. The activity is the same, the Aspect Marker (the Progressive *-ing*) is present in both cases. The time scale of the activity of playing tennis (s_i) is **minutes**. The reference time scale in (1) (s_{R1}) is **hours**, and in (2) s_{R2} is **years**. Hence $s_{R2} \gg s_i$, and $s_{R1} \sim s_i$. Hence, equations (4) - (6) suggest that (1) be interpreted as a single process, and (2) as an extended sequence of several individual tennis playing events. This is true even in combination with temporal adverbials. For example, compare the following two sentences

1. This morning, when McEnroe was playing Tennis, he sprained his ankle.
2. In 1991, when McEnroe was playing Tennis, he sprained his ankle.

Reasoning in multiple time scales is a very useful abstraction to make since it cuts down unnecessary cognitive processing. It also turns out to explain cross-linguistic data. In fact, in languages such as Tamil and Spanish, there is no *Habitual* aspect, and so disambiguation of sentences such as the ones above can only be done using a scheme like the one suggested here.

4 Deriving Verb Classification Schemes From The Primitives

In this section, I derive the canonical verb classification schemes that are a result of previous work in Formal Lexical Semantics. Most previous work on Aspectual Classification of Predicates [Vendler1967], [Comrie1976],[Dowty1979], [Moens1988], [Bennett1986], [Dorr1993] have used the following classes.

1. States.
2. Punctual Activities.
3. Durative Activities.
4. Accomplishments.
5. Achievements.

Though the exact definitions vary across the different proposals, the main distinctions between these classes is made in terms of the following three parameters

1. Telic/Atelic [+t/ -t].
2. Punctual (Atomic)/Durative [+a / -a].
3. Dynamic/Stative [+ d / - d].

The table below shows how the various parameter settings constitute the various classes.

Aspectual Class	**Parameter Setting**
States	[-d]
Punctual Activities	[+a, -t, +d]
Durative Activities	[-a, -t, +d]
Accomplishments	[-a, +t, +d]
Achievements	[+a, +t, +d]

Henceforth, the various classes mentioned above as well as the associated parameter settings will be referred to as the Vendler-Dowty-Taylor (VDT) classification scheme.

In this section, I will attempt to show the following.

1. The classification schemes and the parameter settings mentioned above can be systematically derived from the more process-based parameters described in the last section.
2. Several well known problems in Aspectual Classification (such as Indefinite Plurals, Imperfective Paradox, unclassifiable words (such as die)) are naturally solvable using the process-based parameters mentioned earlier.

In subsequent sections, I will show the overall solution to some problems such as the *Imperfective Paradox* in greater detail. I will also show how the new Aspectual class structure can help in discourse inferences.

Predicate	s_i	**Effort**	**Period**	**Resource**	**Intentional**	**VDT Class**
walk	Minutes	Yes	Multiple	No	Yes	ACT
run	Minutes	Yes	Multiple	No	Yes	ACT
know	Years	No	Aperiodic	No	No (control)	STATE
run a mile	Minutes	Yes	Aperiodic	Yes	Yes	ACC
reach	Seconds	Yes	Aperiodic	Yes	Yes	ACH
push a cart	Minutes	Yes	Multiple	Yes	Yes	ACC and ACT
draw a picture	Minutes	Yes	Aperiodic	Yes	Yes	ACC
believe	?	No	Aperiodic	No	Yes	STATE
die	Days	No	Aperiodic	No	No (- control)	ACH
fall	Seconds	No	Aperiodic	Yes	No (- control)	ACT
climb	Minutes	No	Aperiodic	Yes	Yes	ACT
cough	Second	No	Multiple	No	No	ACT (Point)

Table 1: Classification of some common situations and their Vendler Class.

4.1 Deriving The VDT Scheme

We can now derive the correspondence between the Vendler Dowty Taylor Classification (VDT) system and the process parameters. The table below shows the mapping. An explanation of the mapping scheme follows.

VDT Class	**VDT Parameters**	**Process Features**
States	[-d]	Marginal Effort
Punctual Activities	[+a, -t, +d]	$s_R \gg s_i$
Durative Activities	[-a, -t, +d]	$s_R \sim s_i$
Accomplishments	[-a, +t, +d]	$s_R \sim s_i \wedge$ Resource present
Achievements	[+a, +t, +d]	$s_R \gg s_i \wedge$ Resource present

- **States** in the VDT system correspond to situations which require **marginal effort** to maintain. In many cases, there is some effort required in the situation, as in the the case *sit* or *stand*. In other cases, there may be no effort required as in the case of *know*. The VDT system classifies states as unchanging situations. This creates several anamolies, and in the next section, we show how our definition is a more natural one.

- **Durative Activities** in the VDT system correspond to processes whose time scale (s_i) is on the order of the reference time scale (s_R), ($s_R \sim s_i$). Activities may be effortful (as in *run*, *walk*) or effortless as in *fall*.

- **Punctual Activities** in the VDT system correspond to processes whose time scale (s_i) is **much lesser than** the reference time scale (s_R) ($s_R \gg s_i$).

- **Accomplishments** in the VDT system correspond to processes whose time scale (s_i) is on the order of the reference time scale (s_R) ($s_R \sim s_i$), **and** which requires a specific type of *resource*. This allows for the solution of various problems that arise from the $[+d, +t, -a]$ classification of the VDT system. We consider this in greater detail in the next section.

- **Achievements** in the VDT system correspond to processes whose time scale (s_i) is **much lesser than** the reference time scale (s_R) ($s_R \gg s_i$). Achievements are usually controllable processes.

4.2 Problems With VDT Classification and The Process Solution

We now consider the parameter space used in the VDT classification in greater detail, and see how the process based parameters address the attendant issues.

- **Durative versus Punctual**

  To illustrate the point, consider the famous example below.

  Example 6: Bill is coughing.

  The standard account of the meaning of this sentence is that coughing is a *punctual* verb, and the only way it can appear in a Imperfective perspective (since the Imperfective perspective requires a *Process verb*) is if the process is *iterated.* This is a result of the *inherent* punctuality of cough.

  The argument stated above does not work. For instance, if I change the time scale of reasoning by showing a slow motion movie of a cough and saying *Look he is coughing*, I could be pointing to the process of a *single* cough.

  Instead of the kind of argument made above, I propose that punctuality of verbs is a derived entity that depends on the interaction between a Time Scale of reasoning and the time scale of an activity. In most contexts, since the time scale of a cough is much smaller than the reference time scale, a cough counts as an event as and needs to be iterated to form a process. However, upon changing time scales of interest, a cough can become a process.

  So the classification of verbs as *punctual* and *durative* is **not** an *inherent* attribute of the verb, but a function of the specific context in which the verb is used as well as our experience of processes in the world. Our experience of most activities seem to fall into the **minutes** time scale, and so the *default* interpretation would classify activities in the **second** time scale (like *cough*) as a punctual verb. So while it is *predictable* that some languages may make this default distinction, it is not an inherent property of a verb.

  The Multiple Scale Hypothesis can also elegantly explain when a durative verb can be recast as a punctual verb. We saw an example of this modification in the earlier section with the *play tennis* situation where a shift in the level of the temporal resolution resulted in the shift in the Aspectual Class of the situation.

- **Telic vs. Atelic**

  Telic situations are those that have a well defined end point. For instance *John is singing* is Atelic since it does not have well defined end point beyond which the process cannot continue, but *John is drawing a circle* is Telic since the activity stops when a circle comes into existence. A key test for whether something is telic or not is to test for the Imperfective entailment. For instance *John was drawing a circle* does **not** entail that *John drew a circle*, whereas *John was* *singing* entails *John sang.* This difference leads to the **Imperfective Paradox** [Dowty1979]. We will deal with this in more detail later.

Telic situations are therefore distinguished in standard theories as an **Activity** with a well defined **Achievement**. Hence telic situations are a combination of a **Process** followed by an **Event** which is the culmination of the **Process**.

This notion of telicity creates problems. For instance, take the verb *die.* Vendler [Vendler1967] classifies it as an *Achievement*, which implies a *punctual* characteristic. But the fact that you can say *Harry is dying* suggests that the situation is *durative* and *telic.* So Vendler's classification is clearly problematic.

On the other hand, if we suggest that *dying* is telic, composed of a *process* and then an *event*, we should be able to say ** Harry was dying, but recovered.* Since this is quite a strange sentence for most speakers, [Comrie1976] suggests a new class of situations where *dying* is a *special* telic predicate, one which does not allow its end-event to be prevented, once the associated process starts.

Our model handles this situation quite easily. Let us consider the parameters that characterize the predicate *die*, namely

Parameter	**Value**
Time Scale	Days
Effort	No
Resource	No
Periodicity	Aperiodic, irreversible
Intentional	No, uncontrollable

In this classification, once the process of dying has started, there is no means to prevent it.[4] Contrast this with the parameters for another telic situation such as *drawing a circle.*

Parameter	**Value**
Time Scale	Minutes
Effort	Yes
Resource	Yes
Periodicity	Aperiodic
Intentional	Yes

Here there is clearly a means of ending the process either by not expending effort or by the absence of a resource (drawing material or pencil), or by deciding to stop since the activity is controllable.

In summary, whether a situation is Telic or Atelic is simply determined by the resource requirement and controllabililty of a situation. In our case, this translates to the following:

1. A situation can be considered telic if it requires a Resource. In this case the final event of the characterizing the situation is the complete consumption of the resource or of the creation of a resource. An example is *reading a book*, where the situation ends when the book is read.

[4] Since our model is **strongly operational**, this fact follows from the execution of the model itself. For details of the model, refer to the next section.

2. Any intentional situation can be Telic. In this case the final event signifying completion can be a durative adverbial (as in *He ran until 5 P.M.*), a specific world state (as in *draw a circle*), a specific effort (as in *run till tired*), or a specific *affective* state (*sleep till restored*).

- **Static Versus Dynamic**

  States are classified by their unchanging nature. One of the tests proposed by Dowty is

  Definition : 1

  If α is a stative predicate, than α(x) is *true* at an interval I *just if* $\alpha(x)$ is *true* of every moment in I.

  Clearly this is insufficient. For instance, *sleep* is a stative but *Harry slept for eight hours* is still true if Harry got up a few times to go to the bathroom.

  I propose instead the following definitions.

  Definition : 2

  $\alpha(x)$ is a stative predicate if the **effort** to maintain $\alpha(x)$ is marginal.

  This satisfies all the words that are often referred to as statives (such as *know, believe, sit*, and *sleep*). Some of the statives require marginal effort (such as *sit*, and *stand*).

Having shown how our ontology provides a more natural way to model the Aspectual Class of a predicate, we now on to describe the computational model of Aspect.

5 Computational Model

Aspect is modeled using **Aspectual Processing Nets (APN's)**. Appendix A lists the formal definition of an APN. Formally, APN's are a modified class of Petri Nets [Murata1988], called **Process Nets** where the modifications allow for the representation of stochastic transitions, as well as special compositional states and transitions (decomposable to a an ordered collection of primitive states and transitions).

5.1 Process Nets

Definition : Process Nets

A Process Net consists of *places*(P) and *transitions* (T) connected by weighted directed arcs $\mathcal{A}$ ($\mathcal{A} \in (\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})$). Each arc $a_{ij} \in \mathcal{A}$ has weight w_{ij} (($w_{ij} \in \mathcal{N}$.). Input Arcs $*\mathcal{T}$ ($*\mathcal{T} \in (\mathcal{P} \times \mathcal{T})$) connect Input Places to Transitions. Output Arcs $\mathcal{T}*$ ($\mathcal{T}* \in (\mathcal{T} \times \mathcal{P})$) connect Transitions to Output Places.

The underlying model of the Process net is a weighted, bipartite graph that consists of **Places** and **Transitions**. Arcs are either from Places to Transitions, in which case the Places are *Input Places* of the Transitions, or from Transitions to Places, in which case the Places are *Output Places* of the relevant transition.

Figure 3 shows a very simple net with a single transition $T1$, two Input Places $I1$ and $I2$, and a single Output Place $O1$.

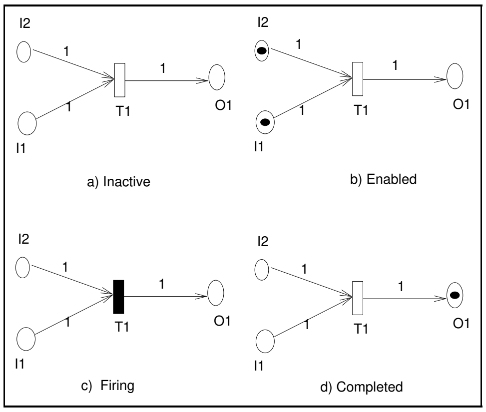

Figure 3: Basic Execution Statuses of the APN

In the graphical representation, Places are drawn as circles, and Transitions as boxes. Arcs are labeled with an *integer* weight. An arc weight of k, can be interpreted as k parallel arcs. Labels for unity weights are usually omitted. A *marking* assigns to each Place a non-negative integer. If a Place p is marked with the integer k, we say that the "*place p* is marked with *k tokens*." Graphically, this is illustrated by the presence of a black dot with the integer k alongside the relevant place p. A marking is algebraically described by an *M-vector*, where the i_{th} element of the vector is $M(i)$, denoting the number of tokens in the i_{th} Place. Clearly, the M-vector ranges over the number of Places in the net.

In order to simulate the dynamic behavior of a system, a Marking of the Process Net is changed according to the following rule, which is the **firing rule** of the transition.

1. A transition t is said to be enabled if each input place p of t is marked with at least w_{pt} tokens, where w_{pt} is the weight of the arc from p to t.
2. The firing of an enabled transition t, removes w_{pt} tokens from each input place p and places w_{pt} tokens in each output place of t.

Figure 4 shows the basic execution statuses of the simple PN.

The various execution stages shown in the figure correspond to the following definitions in the APN.

1. In (a), the net is shown in an *inactive*. An inactive transition corresponds to an unmarked state (the absence of any token in any input place). A net is inactive is all transitions are inactive. Inactive nets are uninteresting.
2. In (b), the net is shown in an *enabled state*. An *enabled transition* (j) corresponds to having at least w_{ij} tokens in each of its Input Places. A net comprising of at least one enabled transition is in an *enabled state*.
3. In (c), the net is shown in a *firing* state. An enabled transition can fire depending on its firing rule and the its type. The following section explains the type of transitions in a PN.
4. In (d), the net is shown in a *completed* state. An transition j that has finished firing deposits w_{jo} tokens in each of its Output places.

5.2 Transition Types

A Process Net consists of three types of transitions, namely *deterministic* transitions, *stochastic* transitions, and *hierarchical* transitions.

Figure 4 shows the various transition types and their graphical representations.

1. Deterministic Transitions fire with a fixed delay after being *enabled*. The actual delay is an integer which varies depending on the nature of the transition. For logical conditions, for instance, the delay is very small, and the transition can be considered instantaneous.
2. Stochastic Transitions are associated with an exponentially distributed random variable that expresses the delay between enabling and firing. Due to the memoryless feature of the exponential transition, PN's with stochastic transitions are isomorphic to Markov Chains [Molloy1982].

Transition Type	*Description*	*Graphical Representation*
Determinstic	Fires after a fixed delay after being enabled.	
Stochastic	Fires after an exponentially dist. delay after being enabled	S
Hierarchical	Activates a subnet and waits for control to return	

Figure 4: A description of Basic Transition Types.

3. Hierarchical Transitions correspond to the *activation* and *execution* of a **subnet**. The semantics of a hierarchical transition is such that the hierarchical transition has to **wait** till the subnet finishes execution and returns control.

The subclass of PN's where all transitions have exactly one incoming and one outgoing arc are called *State Machines*. State machines model *Finite State Machines* (FSA). State machines can model *decisions* (Single Input place to Multiple Transitions), but not the synchronization or coordination of parallel activities.

5.3 Modeling Aspect Using Process Nets

This section details the mapping from our ontology to Process Nets. The mapping consists of two parts.

1. Modeling Semantic Features.
2. The Aspect Processing Net.

We consider the two parts in detail.

5.4 Mapping Lexical Semantics

In this section, we consider the PN model for the semantic features described in Section 3. This section will describe the model for each individual feature that characterizes the Apsectual Class of a situation. Subsequent sections will show how these feature models combine with the overall APN to model the semantics of aspect.

Figure 5 shows the PN model for the individual features discussed in Section 3. These models are described below.

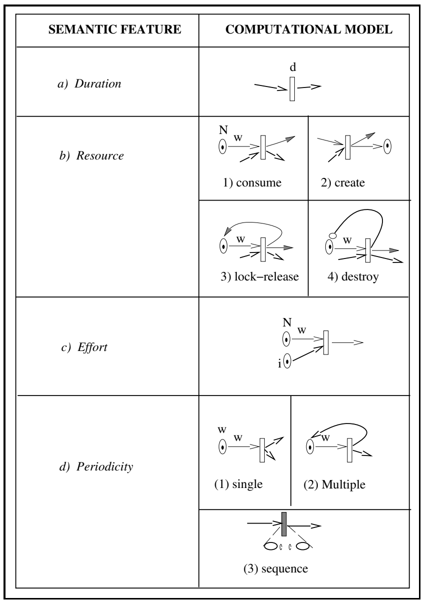

Figure 5: Modeling Lexical Semantic Features.

- **Duration**

  Modeling fixed duration in a process net is trivial. The firing semantics of the transition type models the duration involved. Figure 5 models a situation of duration d.

- **Effort**

  Modeling effort is fairly straightforward. Figure 5 (c) shows the model graphically. The two input places correspond to the the expenditure of energy and intentionality (labelled i). The amount of energy available is the value N shown. The amount of energy required for the situation is w. If $N \gg w$, the situation can be repeated several times.

- **Resources**

  As was discussed earlier, resources can be consumed, created, held and released or destroyed. Figure 5 (b) 1- (b) 4 graphically depict the model of the different types of situations.

  1. The consumable resource has same form in our model as the *Effortful* situation. The distinguishing feature in the two cases is the intentional input which is a requirement for the *Effortful* situation, but not a requirement for the consumable resource.
  2. Creation of a new resource is modeled by the presence of an Output Place that models the created resource. Consider the following example.

     Example 7: John drew a circle.

     The completion of the drew situation results in the creation of a circle (at least a representation of one).
  3. Resources can be locked (if shared) and then released at the end of some situation. Computation resources, such as *attention* can be modeled this way.

4. One interesting thing to note is in (b) 4, where the *inhibitor arc* assures that the resource is unavailable. Consider the following sentence.

   Example 8: The building was demolished.

   In this case, the building is unavailable at the end of the situation. The only way to rectify the situation is to have another process (perhaps a "construct building" process) remark the resource for use.

- **Periodicity**

  Modeling Periodicity is shown in Figure 5 (d). Again the technique is fairly straightforward, with single, multiple and sequential situations shown in the figure. The equation $N >> w$ in the multiple period situation, ensures that the situation can be iterated multiple times. Note there may be other interactions (such as resource or effort) that may ultimately decide if the situation is actually repeated or not in any particular **episode**.

- **Reasoning at Multiple Levels**

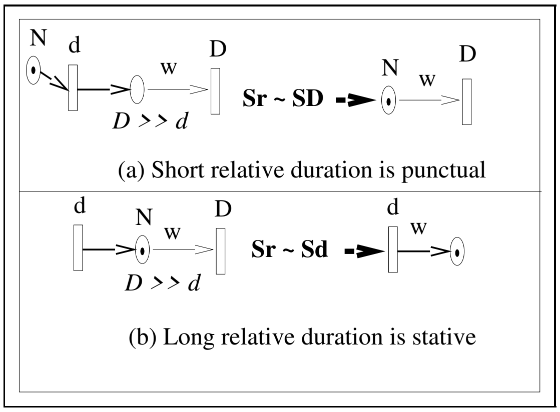

Figure 6: Reasoning At Multiple Scales. An example of dimensionality reduction.

Figure 6 shows the two cases of dimensionality reduction that come from the interaction of a reference time scale (S_r) with a particular situation. In the first case we have the following relationships

$$(S_r \sim S_D) \wedge (S_r \gg S_d). \tag{7}$$

This allows (1) to apply, resulting in the model shown on the right in Figure 6 (a). An example of this situation is when the *enable-start* path of the transition graph in Figure 2 is not represented due to the relatively small interval size of that path.

The other case, namely

$$(S_r \sim S_d) \wedge (S_D \gg S_r). \tag{8}$$

results in the transformation shown in Figure 6 (b). An example of this situation results in the transformation of some situations (such as **know**(ϕ)) into *places* in the transformed PN.

5.5 The Aspect Processing Net (APN)

Now that we have describe the fundamental operations of the Process Net and the model of the individual features, we present the computational model of Aspect, called the Aspectual Processing Net (APN).

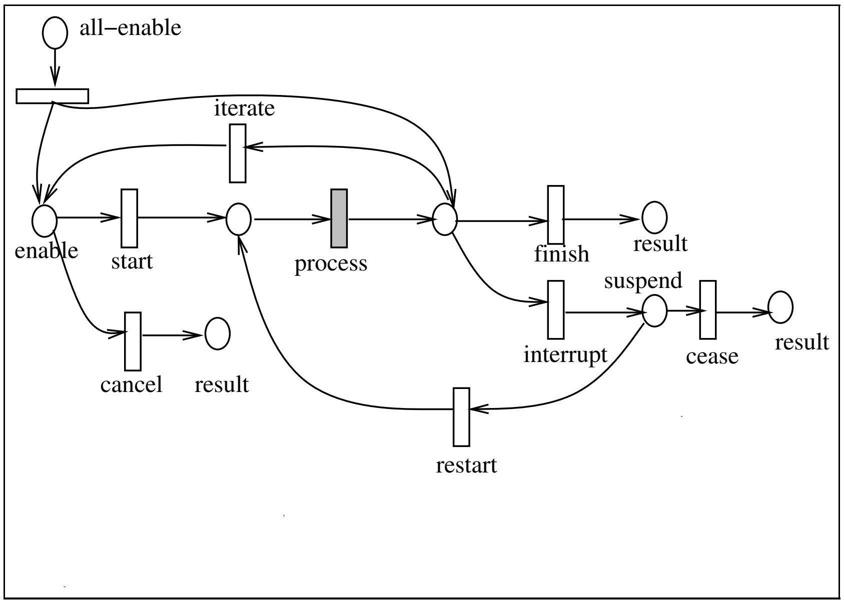

Figure 7: The Aspectual Perspective Net (APN).

The APN described follows the natural ontology of situations shown in the transition graph of Figure 2 (b). The various transitions (*start*, *process*, *finish*, *interrupt*, *cancel*, *iterate*, and *restart*) all refer to links or states in the *transition graph*. There is an additional *all-enable* state that marks both the *enable* and a *process* state. The *process* transition is usually a hierarchical transition in that it is composed of interacting subnets.

In several places of the APN we have multiple outgoing arcs leading to multiple transitions. Such a situation is a referred to as a **decision** or **conflict** point and leads to ambiguity in interpretation. The solution to the **conflict** is provided by the specific combination of the relevant Linguistic Construction (denoting the Aspect Marker System) and the natural semantics of the situation, whose ontology we considered earlier.

To illustrate the disambiguation of a **conflict** consider the Output of the *Process* transition. The Output place could potentially enable three transitions, namely *iterate*, *finish*, and *interrupt*. Let us look at possible selection of transitions.

1. The construction used chooses a specific transition such as the use of *finish* in *John finished playing*.

2. The natural semantics of the situation controls which transition to fire, as in *John read the book*. In this case, the complete consumption of a resource (unread book), controls the firing

of *finish*. Compare this to *John got tired and stopped reading* where the lack of available effort selects the *interrupt-suspend-cease-result* path.

3. The Aspect Marker system may prefer specific paths, as in the case of the *used to* construction. In this case, the selected transition is *iterate*.

Grammatical Constructions, Natural Semantics and other ways of expressing aspect interact with the APN to jointly produce an interpretation. The general mechanism for this interaction is outlined below.

- **Grammatical Constructions** (the Aspect Marker system) provide an *initial marking* to the APN. The initial marking may consist of placing tokens in selected place(s) of the APN enabling one or more transitions. The Aspect Marker system may also establish a specific resource to preferentially select a given transition in case of a **conflict**.
- **Other Ways Of Expressing Aspect** may place one or more tokens in specific parts of the APN. This may enable one or more transitions in the APN.
- **Natural Semantics** of situations are modeled using the *features* described earlier. The PN's corresponding to the specific characteristic features for a situation interact with the APN by enabling or disabling specific transitions. The *process* transition is usually decomposed to the appropriate situation specific subnet.

The next section looks at modeling the various features of the situation semantics in greater detail.

5.6 Grammatical Markings

Linguistic Constructions (the Aspect Marker System) provide an *initial marking* to the APN. As discussed earlier, an initial marking of the APN corresponds to placing *tokens* at various places on the net.

Figure 8 illustrates the case with the English *used to* construction. Here the *iterate* transition can be enabled. Notice, that *conflicts* are still unresolved, since the natural semantics can still interact in the following ways:

1. The used to construction can interact with other aspects (such as the *progressive* in *Harry used to be playing*). This kind of interaction is considered in section 7.
2. The natural semantics *block* or *modify* the *iterate* transition. Consider the example **Harry used to be dying.* Clearly, the irreversible, and aperiodic nature of the predicate *die* blocks the *iterate* transition in this case.

The number N shown beside the Input Place to the iterate transition, signifies the integer number of tokens. This can be unbounded, as is the usual case, or bounded as in *McEnroe used to practice everyday till he was 25.*

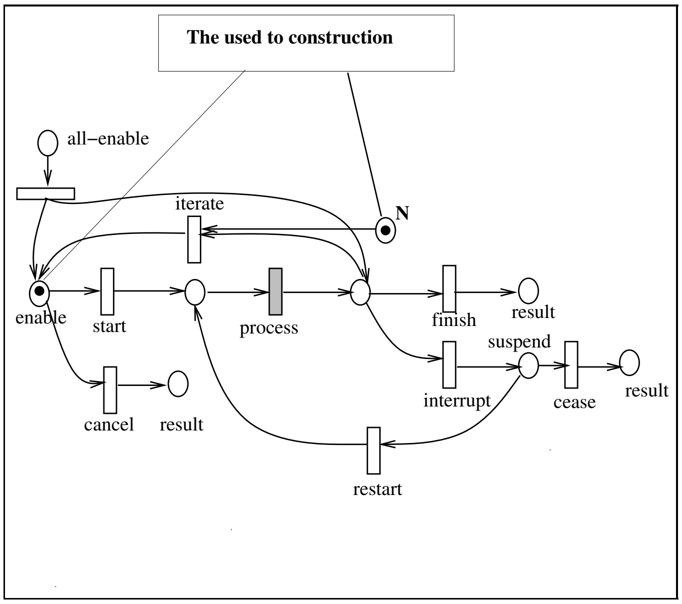

Figure 8: Constructions Mark the APN : Example Of the *used to* Construction.

Other Constructions that pertain to the Aspect Marker system of the world's languages supply different markings to the APN. The details of the specific markings form the subject of discussion of the following three sections.

5.7 Expressing Aspect

While most languages have grammatically obligated structures for expressing aspect, there are many optional devices used as well. In Section 2, we discussed some extended aspectual forms in different languages. The table below shows the various models of the extended forms in the APN.

Aspectual Expression	**Modeling with APN**
Cessive	Enable cease
Completive	Enable finish
Cumulative	Multiple tokens on enable
Imminent	Mark enable
Inceptive/Inchoative	Enable start
Iterative	Enable iterate with multiple tokens
Protracted	Enable process with multiple tokens

Figure 9 graphically illustrates our model for the **cessive** (ex. *stop*), **cumulative** (ex. *one by one*), and **Protracted** (ex. *keep on, continue*) aspects.

The relevant part of the APN that is marked is shown in Figure 9. The mapping to the other forms of extended aspect are straightforwardly translated from the table above. The following observations can be made about the model.

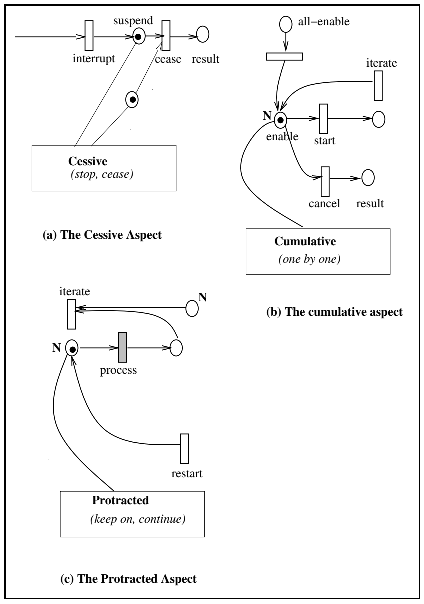

Figure 9: Extended Aspect: Example Of the (a) cessive, (b) cumulative, and (c) protracted aspects.

1. In the case of the **cessive** form, we see that the *cease* transition is enabled. This allows all inferences that follow from the resulting stage of the cease transition automatically. For instance, in the interpretation of the sentence *John stopped exercising*, the automatic inference *John stopped expending effort* would be available upon inspection of the APN.

2. In the case of the *cumulative* aspect (and any other instance of expressing countable instances of a situation), the placing of multiple tokens allows for the same transition to fire multiple times. For instance, to model the sentence *Five soldiers crossed the border*, the *start* transition of the *cross-the-border* situation is enabled with 5 sets of tokens. [5]

3. In some cases, the supply of tokens to a place may be unlimited, as in the case of a *protracted* situation. For example, in the case of *keep on playing*, the *process* transition characterizing the *playing* situation is continuously enabled. Other constructions such as *until* (as in *keep on playing until you tire*) may inhibit the *process* transition when a specific condition (tire) obtains.

5.8 Lexical Modification

Lexical Modification of the aspectual perspective pertains to the interaction of the semantics of the situation with the APN. An example shown in Figure 10. The *Perfect* perspective that results from the *has fallen* construction is modified by the semantics of falling into a hole. Let us consider the various features of the falling that comprise the semantics of falling into a hole situation.

[5]Note the tokens can be typed (colored) to model discrete entities such as *The general and two soldiers crossed the border.*

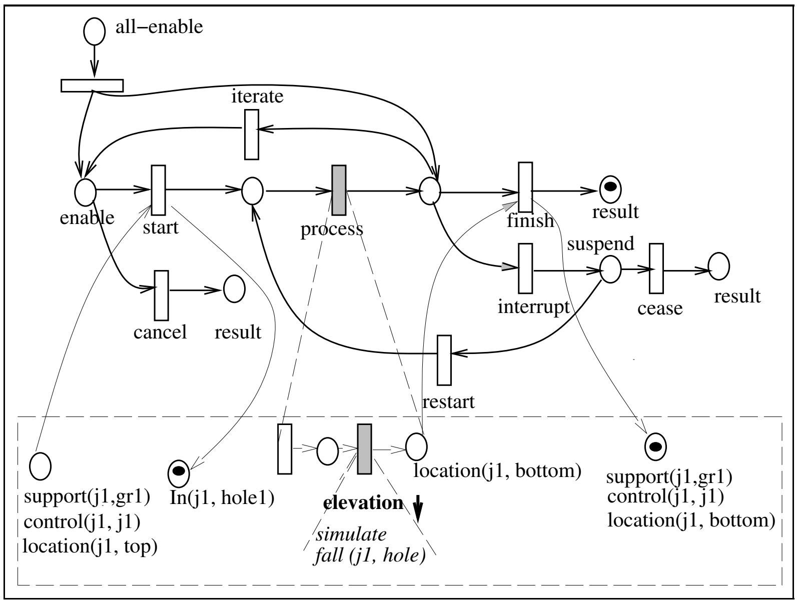

Figure 10: Lexical Semantics Modify APN: Representation of *John has fallen into a Hole.*

The various relevant parts of the model are explained below.

1. At the beginning of the situation, the required preconditions are that John be supported, in control of his location, and be at the top of the hole. This is modeled in the current system by the attributes (support(j1, gr1) $\wedge$ control(j1, gr1) $\wedge$ location(j1, top).[6] The presence of a token at the place representing these conditions is a prerequisite to starting the fall transition. In our model this is simply indicated by the arc from this state to the *start* transition.

2. Once the fall episode starts (a result of firing the *start* transition), John is unsupported and does not control his location. This is **implicitly** represented by the *absence* of a token at the place described in (1). Inspecting the APN directly yields this inference. This is part of the **Strong Operationality** principle.

3. The completion of the *start* transition also results in a marking on the condition *in(John1, hole1)*. Note that this condition is active throughout the **episode** of reasoning, again allowing for the automatic inference at any point that *John is in the hole.*

4. The process of falling is a hierarchical transition that activates a subnet that simulates the fall (elevation decreasing, etc.). Again, the specific reference time scale controls the granularity of reasoning.

[6]In general, each condition is a separate place, and each argument (support (x, y)) is *typed* using a specific color for a type (a color for x and one for y). Without loss of generality, we will omit the colors, and collapse the three conditions to a single place in our graphical representation. A net with a finite number of colors can always be *unfolded* into a single colored net.

5. The *perfect* nature of the *has fallen* construction (which is described in greater detail in a later section) allows for the resulting state of the fall available for further processing. Again, this is a part of the execution of our model, in that the all the conditions that obtain as a result of falling into a hole ((support(j1,gr1) $\wedge$ control (j1, j1) $\wedge$ location(j1, bottom))) are marked and available for processing.

Now that we have described the fundamental operations of the active computational model of aspect, we turn to some detailed issues in modeling the Aspect Marker System of the world's languages.

6 Imperfective

An *Imperfective* aspect involves viewing a situation with explicit reference to its internal consistency. The internal consistency of a situation includes the type and duration of the phasal sequences that comprise the situation.

Example 9: John is speaking.

In this case the *process* transition is enabled which automatically allows the inference of the ongoing nature of the speaking event.

Many languages have a distinct category to express imperfectivity. Others have more specific categories that denote partial groupings of the meaning of imperfectivity. Figure shows the various relevant groupings. The following table shows some examples of languages and their use of *separate categories* for these aspectual groupings [7]

Language	**Imperfective**	**Habitual**	**Continuous**
Spanish	Yes	No	Optional
English	No	Only in the past (*used to*)	Progressive
Tamil	Yes	No	No
Russian	Yes	No	No
Hindi	No	Yes (*-ta*)	Progressive (*raha*)

From the table above, we see that while Russian and Tamil grammars have a single form for the Imperfective that covers both Habitual and Continuous aspects (including progressive and non-progressive forms), English has a special habitual form (*used to*) in the past tense, and a progressive (morphologically indicated by *-ing*). Spanish, on the other hand has an optional Progressive form though the Imperfective form does not exclude a Progressive reading. Hindi has a special habitual form (morphologically marked by the suffix *ta*) as in *Ladka yahan kam karta ta* (The boy used to work here). Hindi also has a separate construction for the continuous aspect *raha* as in *Ladka yahan kam kar raha ta* (The boy was working here). A literal translation of the previous example is *The boy (Ladka) here (yahan) work (kam) make (kar) inhabit (raha) was (ta).* The literal translation of *rahna* is inhabit, showing the connection between aspectual marking and location. This connection is explored further in Section 11.

The APN can simulate the different groupings found in the different languages quite easily. In our model, we have the following mappings.

1. Imperfective Aspect corresponds to marking the **process** place and the preferentially selecting the *iterate* transition. Languages that have a separate category that covers the entire imperfective use this mapping. Figure 11 shows this process.

2. Progressive Aspect corresponds to enabling the **process** transition of the Aspectual Perspective Net. Languages which have a separate category to express the Progressive aspect make use of this mapping. Figure 12 shows this mapping graphically.

[7] The Russian, and Spanish distinctions are from [Comrie1976].

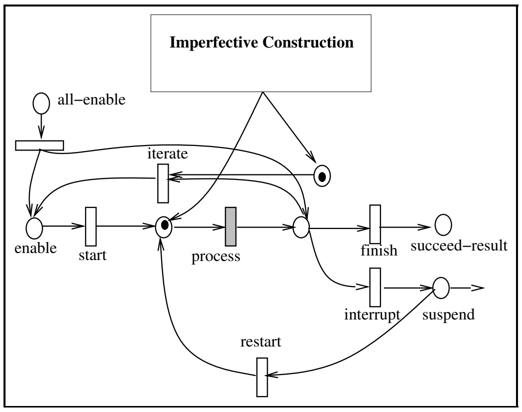

Figure 11: Modeling The Imperfective Aspect: The Imperfective Marking on the APN

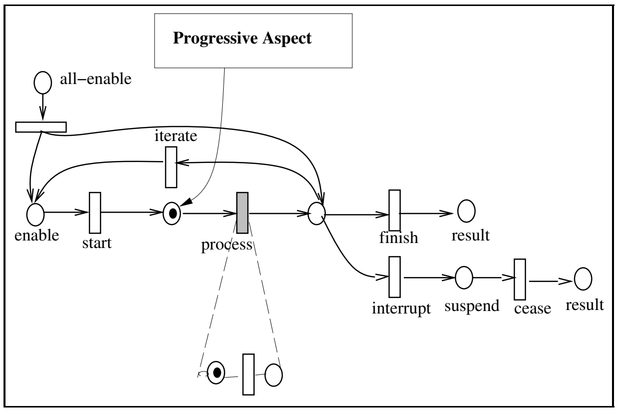

Figure 12: Modeling The Progressive Aspect: The Progressive Marking on the APN

3. Grammatically marked Habitual aspect maps to preferentially selecting the *iterate* transition. Notice in English, this corresponds to mapping the explicit use of *used to.* Figure shows the Habitual Perspective in the interpretation of the sentence *John used to work here.* In Hindi, this corresponds to the form *raha.* Figure 13 graphically shows this model.

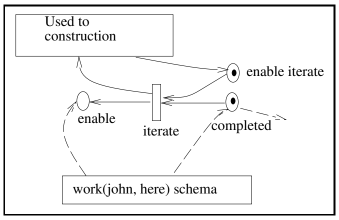

Figure 13: The Habitual Perspective with the *Used to* Construction.

4. Certain languages such as Spanish allow both an imperfective as well as a separate Progressive form (or Habitual) use the appropriate combination of the earlier techniques to denote aspectual perspective.

[Comrie1976] lists some additional issues with the Habitual and Progressive aspect that are considered below.

6.1 Habitual and Other Aspects

Since most situations can be extended over a period of time or iterated sufficiently many times, we would expect most languages to allow for combination of the Habitual with other kinds of semantic aspectual values. Indeed, most languages allow for the combination of Habitual with Progressive as *used to be working* in English or *kam karta rahta ta* (used to be working) in Hindi. Also common is the combination of the Non-Progressive with Habitual as in *He used to work here.* In the case where Habitual expresses iterativity, many languages allow the use of the Habitual with the Perfective where an individual situation is expressed in the Perfective, and the iterated version of the situation is expressed in the Habitual sense. Many languages allow for the use of special adverbials such as the Russian *byvalo* (it used to be), or the Hindi *hotha tha* (it used to happen) to establish habituality, and then allowing for the combination with the perfective to describe each individual occurrence of the habitual.

The APN model explains each of these combinations. the mappings for each of these categories.

Combination Type	**Modeling with the APN**
Habitual + Progressive	Enable *process* and *iterate*
Habitual + Non-Progressive	Mark *enable* and *iterate*
Habitual + Perfective	Enable *iterate* and mark *all-enable*

The following automatic inferences can be obtained from the executing APN.

1. In the Habitual + Progressive case, such as in the example *John used to be singing*, the APN automatically allows for the inference that the process of singing had already started at the

time of interest. For instance is I said *John used to be singing when the sun set*, inspection of the APN would allow us to conclude that the singing began prior to the sunset.

2. In the Habitual + Non-progressive case, such as in the example *John used to sing when the sun set*, the APN allows for the automatic inference that the *start* of the singing *followed* the sun set, since the start transition would fire only *after* the sun set.

6.2 The Imperfective Paradox

The Imperfective Paradox [Dowty1979] comes from trying to separate verbs of accomplishment from verbs of activity. One diagnostic test comes from the different entailments of the two verb classes when used in the Progressive Aspect. Dowty illustrates the problem with the following set of examples

1. John was drawing a circle.
2. John drew a circle.
3. John was pushing a cart.
4. John pushed a cart.

Here (1) does not entail (2), but (3) entails (4).

However, *drawing a circle* involves bringing a circle into existence, and yet it is this entailment that fails when the sentence appears in the Progressive Aspect. This is called the *Imperfective Paradox.* The solution of the problem requires an explanation of how *drawing a circle* entails the activity of bringing a circle into existence but **does not entail** that a circle was bought into existence.

In a solution to the problem, Dowty proposes that the truth conditions for the Progressive aspect depends on a new primitive called *Inertia Worlds. Inertial worlds* are the set of predictable future worlds. [8] In the *drawing a circle* case there is a predictable future world in which the existence of the circle (at least a representation of a circle) is true.

While the use of time (intervals) and the requirement of a model theoretic semantics forces Dowty to have to invent the external primitive called *Inertial Worlds*, the APN model solves the Imperfective paradox implicitly. The solution is in the Operational Semantics of the active APN model itself.

The relevant part of the operation of the APN that solves the Imperfective Paradox of (1) and (2) is shown in Figure 14. Summarizing the operations shown in the figure, we have the following:

Illustrating this in the *drawing a circle* example, the *finish* transition fires *iff* there is a circle resource that is created. Hence the entailment *drew a circle* is only true if the final transition were to fire.

However, the process may be interrupted and the drawing process suspended as in

[8] For example, in the case of a coin toss, the set of inertial worlds would include both the one in which the coin came up heads, and the one in which it came out to be tails.

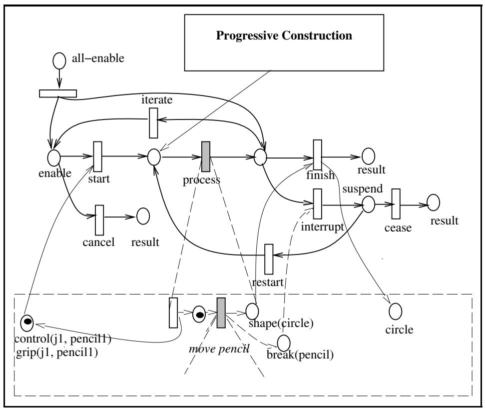

Figure 14: Solution to the Imperfective Paradox: John is drawing a circle.

Example 10: John was drawing a circle when his pencil broke.

This way the token may be held by the **suspend** place, thereby *never* making the existence of the circle true. The dashed lines in Figure 14 lead to the **episode** *drawing-circle process - break(pencil) - interrupt - suspend* which is the desired interpretation of Example 10. Note the absence of a token in the place that represents the circle implies the absence of the circle in the world described by Example 10. This is available by inspection of the executing model.

In contrast, pushing a cart only requires *effort*. Hence the process, once started, entails the the fact that the cart was pushed, since all the transitions (reachability graph) leading from the output place of the **process** transition allow for the interpretation.

To generalize from the example, the solution of the Imperfective Paradox is implicitly represented in the APN model in the following way.

1. The Progressive is modeled by the initial marking that enables the **process** transition.

2. At the stage focussed on by the Progressive, the **process** transition fires by removing a token which will only appear at the resultant state upon completion of the process which requires a specific resource state (creation or destruction).

3. Since at the stage of focus that constitutes the Progressive, the token is absent from the result state, the corresponding entailment of the resultant state is absent as well.

4. If the situation referred to has a specific time scale (say **mins**) as in *I am running two miles today*, and if the next sentence refers to a time much greater than the situation time scale ($\gg$ **mins**), the situation will be assumed to have completed. So if the next sentence is *Tomorrow, I will try three*, the APN model will allow the *automatic default inference I ran two miles.*[9] Note that the situation can always be suspended explicitly with the use of a temporal adverbial like *when* as we saw earlier.

[9] No model proposed earlier (including [Dowty1979] or [Moens1988]) handles this kind of inference automatically.

Thus the Imperfective Paradox is solved implicitly without resort to any external primitives such as *Inertia Worlds*. Also, our solution is part of the model itself, and does not require special type coercion (point to process) as in the case of Moens and Steedman [Moens1988].

7 Perfective

Perfective aspect involves viewing a situation as a single whole, without explicit reference to the individual constituent phases. English does not have a Perfective/Imperfective morphological distinction, but several languages such as Russian, Mandarin, Spanish, etc. do. [Comrie1976].

An example from Russian is the use of the prefixes *po or pos* as in

Example 11: On *po*stojal (*Perfective*) tam cas (He stood there for an hour).

Contrast this with the Imperfective

Example 12: On stojal tam cas (He stood (was standing) there for an hour).

In English, a perfective perspective is involved in the sentence,

Example 13: He fell, got up, and climbed out.

Here each individual situation is viewed as a whole, without reference to the internal temporal consistency.

In our model, a Perfective perspective is modeled by placing a token in the place labeled *all-enable*. Placing the situational token here enables the internal transitions simultaneously. This has several consequences which are explained below.

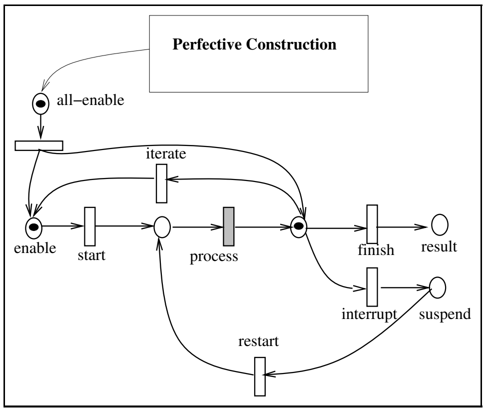

Figure 15: Modeling The Perfective Aspect: The Perfective Marking on the APN.

Since all phases are enabled simultaneously, perfective aspect differs from merely completed action. In our model, aspect refers to a *complete* as opposed to *completed* action, in that the focus is on the entire situation as opposed to the terminal state. This succinctly explains some of the special cases identified in [Comrie1976],such as the presence of the Perfective Future in Russian, as in *Ja ubju tebja* (I shall kill you), by simply allowing the reference to the future event of kill in its entirety.

[Comrie1976] indicates another test specifically for Russian, that imperfective (such as *stoja* (stand)) have a future tense that is different from a present tense (*Ja stoja* (I stand)) and that such verbs can be formed with the auxiliary *budu* (shall be) as in *Ja budu stojat* (I shall be standing). Perfective verbs have only one non-past tense (usually future) and cannot be used with *budu*. Again, this is also succinctly modeled in our approach since our model of the perfective does not allow a token to be exclusively enable the *process* transition (which is needed if we are to have a "I shall be standing" aspectual perspective).

8 Perfect

Perfect Aspect is different from the Perfective/Imperfective distinction discussed earlier since it pertains to the *continuing relevance* of a past situation. Perfect thus is a grammatical way to *relate* situations, and hence provides an inter-situational aspectual perspective. Consider the example depicted in Figure 10:

Example 14: John has fallen into a hole.

Compared to the non-perfect version *John fell into a hole*, we see that the perfect (has fallen) seems to imply that John is still in the hole, whereas the simple past (fell) does not carry this implication. In fact, if John is not still in the hole the use of *has fallen* would be strange and infelicitous.

Note that the perfect need not be confined to relating a past to a present situation. Both the described and the reference situation could be in the past, as in the example *John had fallen into a hole*. In this case, the both the situation of interest and the described situation are in the past. This specific case also has a special name and is called the *Pluperfect* or *past perfect*.

8.1 Modeling Perfect

Each APN models the aspectual perspective that corresponds to a single situation. When there are two situations being related to each other, as is the case with the Perfect, we need two APN's that are related to each other. The two situations correspond to the *Described* situation and the *Reference* situation, and correspondingly we label the two APN's as the DAPN and RAPN.

Notice that in our model the only requirement is that the the result of the DAPN be available for further processing. For example in the case of the sentence *John has fallen into a hole*, we expect the current location of John to be at inside the hole. This is illustrated in Figure 10.

In fact, all the consequent effects of John falling into a hole (location, goal of getting out, etc.) are assumed to hold in the beginning of the reference situations since markings automatically *persist* in places of the APN.

8.2 Specific Types Of Perfect

[Comrie1976] lists some specific manifestations of the the Perfect. We now consider these manifestations in detail. The current use of DAPN and RAPN easily model most of the different types of Perfect. In one case, special modeling primitives have to be introduced. The issues are discussed below.

- **Perfect Of Result.**

  In the perfect of result, a present state is a result of a past situation described by the Perfect. In answer to the question *Is John in the hole now? Yes, he has fallen in*, is a more reasonable answer than *Yes, he fell in*. Similarly, we expect that in the case of the Perfect as in *John has fallen into a hole*, we expect the resultant conditions (such as John's location at the bottom of the hole) to persist in the current state. In other languages, such as Tamil, the the semantics of the Perfect can be decomposed indicate the persistence of result as in *Kannan vilu-in-thi-irukkan*, which literally translates to *Kannan fall-down-state-is*, indicating that the current state of Kannan in the consequent state of fall down. Perfect of result is handled directly by the APN.

- **Perfect Of Experience**

  The Experiential Perfect holds only that a given situation has occurred at least once in the past. There is no necessity that the result of the situation hold in the present.

  In English the semantic difference between *go* and *been* in *John has gone to France* versus *John has been to France* makes this distinction. In the former case, we could conclude that John is still in France, whereas in the latter case all we can say that there was at least one occasion in which John was in France. Other languages, such as Mandarin also make this distinction.

  In our model, the experiential perfect translates to the time ordering of the APN's satisfying:

  $$time(DAPN) < time(RAPN) \leq CurrentTime. \tag{9}$$

  In English, the specific use of constructions such as *since* as in *John has been to Iraq since the Gulf war*, add an additional lower bound to the equation. In order to model this, we need a method to control the **interval of persistence** of tokens. The control of token persistence is not modeled in the effort described here.

- **Perfect Of Recent Past**

  In English, constructions such as *recently* or its synonyms as in *I have recently learned that John fell* can be used in the Perfect form merely by virtue of temporal closeness. The actual amount of time where this relevance is valid depends on pragmatic factors and also varies across languages. For instance in English, *I have eaten this morning* cannot be used in the afternoon, but in other languages (such as Spanish) this may be allowed. Note however, I can say *I have eaten today* since the present moment is included in the definition of today.

  Special constructions such as *recently* mark the consequent state of the described situation to achieve the Perfect effect. The different linguistic conventions that allow for different recency intervals to be used with the Perfect aspect have to modeled as extra constraints outside the APN and are not directly considered here.

  The way to model this would be to have some **expected persistence interval** (possibly related to reference time scale) of tokens in the APN. So the result will be expected hold for a certain time interval before the marking is changed. The exact nature of this model is yet to be determined.

  Once modeled, the semantics should also correctly predict that the while recency implies continuing relevance, the converse is not true. So *The cold war has ended* or even in certain cases *The second world war has ended*, are acceptable sentences.

9 Inferences in Narrative Interpretation

10

In this section, we will consider how our model of interconnected APN's can result in different types of inferences in narratives. We will restrict our discussion to simple narratives, where simple is defined as narratives having a single *discourse segment* ([GroszSid1986]). This basically means that we will consider narratives that deal with a single topic, with a single narrative perspective.

In dealing with modeling the semantics of narratives, we focus on how the choice of aspect constrains the set of allowable inferences. We also show how our choice of an active model of interconnected APN's provides a computational model of automatic inferences. We have already seen how the use of the Perfect establishing certain connections between two situation, thereby constraining the kinds of inferences allowed. We now consider some more detailed examples where our model automatically makes certain commonsense predictions which other theories of semantics have to resort of some metalevel assertions, or explicit rules.

9.1 Chaining

- **Causal Chain** Consider the following narrative.

  Example 15: John fell into a hole. John started to climb out.

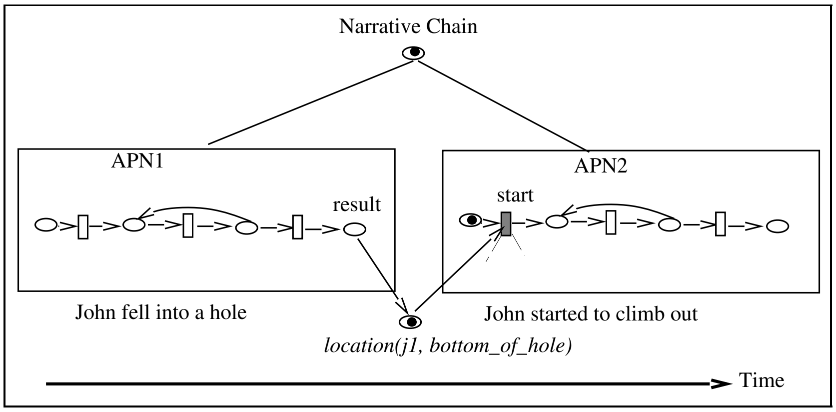

Figure 16: A Causal Inference

- **Resource Enabling**

  Consider the narrative:

  Example 16: The teacher entered the classroom. A student was waiting to see him.

[10]This section is still to be developed

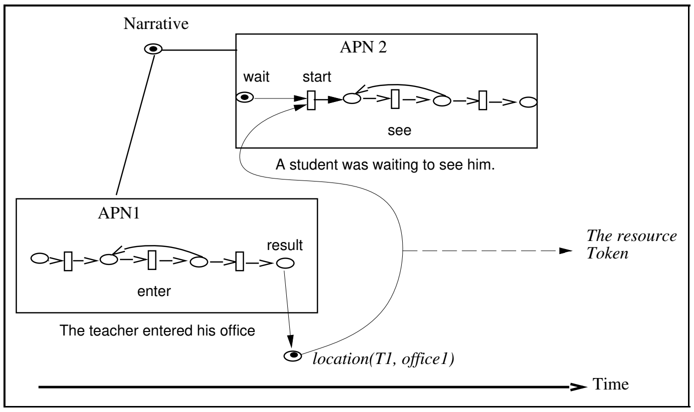

Figure 17: A Forward Bridging Inference

9.2 Aggregation

To be developed

9.3 Missing Inferences

To be developed

9.4 Scripts

To be developed

10 Discussion

11

10.1 Locative Expressions and Aspectual Oppositions

In many languages, the *progressive* aspect is referred to by locative constructions as in *in the process of*, or the Hindi *kar(make) raha (inhabit) ta (was)*. This parallelism between locative expressions and aspect is implicitly modeled in the APN model, since enabling a particular perspective is done by providing a *marking* to a *place*. Our model of aspect also implicitly models a part of he Event Structure Metaphor [Lakoff1992], since the following correspondence holds

Entity	**Event Structure**	**Model**
States	Locations	Marked place
Actions	Motion between locations	Changing markings between places
Goals	Marked Locations	Marked Place

If we superimpose a **scanning** direction of *left-to-right* on the APN, then we get the following additional correspondences between the APN and the Event Structure Metaphor.

Aspect	**Event Structure**	**Model**
Progressive	Moving Forward	Left to right motion of token
Imperfective	Location inside a container	Token in a subnet
Habitual	Visit location multiple times	Mark multiple times
Perfective	View from outside	Net collapsed into single place

Yet to be developed

10.2 Neurophysiological Evidence For Process Semantics

10.3 Connectionist Implementation

- **Using Gradient Descent Over Potential Fields to Implement the APN.**
- **Implementing Scanning.**
- **Binding Token Colors to Phases.**

10.4 Extended Narrative Inferences

To be developed

[11]This section is still to be developed.

11 Conclusion

In the introduction, certain claims were made about the APN model. The APN model described in this report attempted to demonstrate these claims. Specifically, the following novel facts characterize the proposed model.

1. A new set of parameters were proposed to model Aspectual Form of a language as well the Aspectual Class of situations. These parameters suggest a natural ontology for situations based on the logic of processes.

2. An active model was proposed to model aspect. The model satisfies the **Strong Operationality principle** in that the semantics of aspect, including possible entailments and inferences can be obtained by inspecting the executing model itself.

3. The model of Aspect proposed can model cross-linguistic data.

4. The computational model proposed here appears to be capable of addressing broader issues in modeling narrative inferences. More work needs to be done to demonstrate the feasibility of this claim.

References

[Comrie1976] COMRIE, BERNARD, 1976. *Aspect* Cambridge University Press.

[Dowty1979] DOWTY, DAVID, 1979. *Word Meaning and Montague Grammar.* Synthese Language Library, Reidell, New York.

[Moens1988] MOENS, MARC AND STEEDMAN, MARK. 1988. Temporal Ontology and Temporal Reference In *Proc. ACL–88, V4, Number 2, June 1988, pp. 15-29.*

[Murata1988] MURATA, TADAO. 1989. Petri Nets: Properties, Analysis, and Applications. In *Proc. IEEE–89, V77, Number 4, April 1989, pp. 541-576.*

[Nakhimovsky1988] NAKHIMOVSKY, ALEXANDER. 1988. Aspect, Aspectual Class, and the Temporal Structure Of Narrative. In *Proc. ACL–88, V4, Number 2, June 1988, pp. 29-44.*

[Webber1988] WEBBER, BONNIE. 1988. Tense as Discourse Anaphor. In *Proc. ACL–88, V4, Number 2, June 1988, pp. 61-74.*

[Mcdermott1982] MCDERMOTT, DREW. 1982. A Temporal Logic For Reasoning About Processes and Plans In *Cognitive Science 6: pp. 101-155* .

[Vendler1967] VENDLER, ZENO, 1967. *Linguistics in Philosophy* Cornell University Press, Ithica, New York.

[Hobbs1985] HOBBS, JERRY AND MOORE, BOB, 1985. *Formal Theories Of The Commonsense World* Ablex Publishing Corporation, NJ.

[Talmy1985] TALMY, LEN. 1985. Lexicalization Patterns:semantic structure in lexical forms In *Language Typology and Syntactic Description, V 3, pp. 59-138* Cambridge University Press.

[Dorr1993] DORR, DREW. 1993. Interlingual Machine Translation In *AI Journal, V63 (1993) : pp. 429-492* .

[Bennett1986] BENNETT, P.A. ET AL. 1986. *Multilingual Aspects Of Information Technology* Gower, Brookfield, VT, 1986.

[Molloy1982] MOLLOY, M.K., ET AL. 1982. Performance Analysis Using Stochastic Petri Nets *IEEE Transactions on Computers, C-31, No.9, pp.913-917* .

[Berman1994] BERMAN, RUTH, AND SLOBIN, DAN, 1994. *Temporal Inferences in Narrative Interpretation.* ??.

[GroszSid1986] GROSZ, BARBARA, AND SIDNER CANDACE. 1986. Attentions, Intentions, and the Structure of Discourse. In *Computational Linguistics, 12(3), pp. 175-204* .

[Lakoff1992] LAKOFF, GEORGE. 1992. What is Metaphor? In *Advances in connectionist and neural network theory, V2: Analogical Connections,* . Ablex.