Contextual Vocabulary Acquisition

Contextual information in verb contexts: from analysis to algorithm

Chris Becker

Abstract

My goal for this project was to research the verb acquisition process and determine how the CVA verb algorithm could be further enhanced to best define an unknown verb. For part of this research I reviewed the work of some other researchers who have developed their own systems that perform contextual vocabulary acquisition, as well as research relating to the types of contextual information we use to define an unknown verb from context. Next I describe an experiment I did to identify the types of information within a context and how they help to define a particular verb in it. Finally, I describe a computational implementation of a context using some of the inferences drawn from my research and analyze the results.

1 Introduction

Contextual vocabulary acquisition is defined as the 'active, deliberate acquisition of word meanings from text by reasoning from contextual clues, prior knowledge, language knowledge, and hypotheses developed from prior encounters with the word, but without external sources of help such as dictionaries or people' (Rapaport & Kibby 2002). The goal of the CVA project is twofold: to develop and implement algorithms to define an unknown noun, verb, or adjective from a semantic representation of a sentence, and to develop and curriculum to enhance vocabulary learning. The research done by each half is then cycled back to the other to further the development of each goals.

The research on the computational portion currently involves testing and enhancing the existing algorithms for nouns, verbs, and adjectives. Testing is done by representing the semantics of a passage containing some designated 'unknown' word in SNePS (Semantic Network Processing System) (Shapiro & Rapaport 1987). The strength of SNePS lies in its reasoning abilities, and as such provides a powerful framework on which to implement various theories of CVA. A representation of a passage in SNePS will consist of three things: A set of background knowledge containing common-sense facts relating to the 'known' concepts in the passage, a set of background rules, embodying a set of common-sense inferences possessed by a typical reader of the passage, and the semantic representation of the passage itself. With the forward inferencing capabilities of SNePS, the background knowledge and rules will be applied to the information from the passage, generating a set of newly inferred information. The appropriate algorithm will then be executed and will search the entire network for the relevant information needed to construct a possible definition for the unknown word.

2 Previous CVA research

The field of contextual vocabulary acquisition contains a large disparate body of literature from the point of view of various disciplines. In this part of the paper I will highlight some of the work that has provided a foundation for the CVA project. The actual body of literature that the CVA project is based off of is much too large to summarize in one paper. As such I will focus on work that relates closely with my own research for this project.

2.1 Similar Systems

Over the past few decades a number of similar approaches have been developed for computational CVA. Each of these systems has strengths, which we are hoping to mirror, and weakness which we are hoping to improve upon.

2.1.1 Granger (1977)

Richard H. Granger's system, called Foul-Up, is probably the earliest system designed to use contextual information to define an unknown word. The system works in conjunction with scripts, which represent stereotypical, domain specific information which serve as a set of expectations for a particular context. When it comes across an unknown word, Foul-Up will attempt place it in a specific slot in the script representing a concept that the unknown word may fall under.

One of the methods Foul-Up employs to categorize an unknown verb is to categorize it into one of the 10 primitive categories posited by conceptual dependency (Shank & Reiger 1974). The system uses class inclusions of the arguments as well as prepositions to match the verb to a specific primitive category. Using this information, the system can then match the unknown word to a known concept in the script.

Although the system is extremely limited by its reliance on domain-specific scripts, there are some elements that we have found to be worth borrowing for our own system, namely, verb categorization based on prepositions and argument classes. The only drawback is that conceptual dependency primitive categories do not provide a sufficient level of detail to be meaningful in an output.

2.1.2 Carbonell (1979)

Jaime G. Carbonell's system, POLITICS, works in a similar fashion to Foul-Up in that it utilizes conceptual dependency as its representation framework. The system is able to learn the meanings of words by their semantic and syntactic constraints. Using this information it produces a set of expectations for the kinds of contexts a word can occur in. If a conflict occurs between the input and the expectation, then a form of belief revision takes place in which conflicting a rule or fact is removed.

The limitation of POLITICS is that it 'can only induce the meanings of concrete proper nouns', and only when there is 'sufficient contextual information available.'

2.1.3 Hastings (1998)

Peter Wiemer-Hastings' system, Camille, is capable of inferring verb meanings from context using a much larger body of clues than any of the precursors to this system. Hastings' system functions by attempting to place an unknown concept into a predefined ontology.

Many of the ideas that Hastings posits have already been incorporated into our own algorithm. These include the following:

Determining the syntactic frame (i.e. transitive, intransitive, ditransitive).
Identifying reflexive construction.
Determining that there is a syntactic subject (agent).
Determining that there is a syntactic object.

Determining that there is an indirect object.
Determining the actor and object's semantic categories.
Identifying 'with' construction (i.e. identifying instruments)

2.2 Previous CVA research

In addition to computational CVA, our work is built on numerous other research projects on the psychological aspects of CVA. Much research has been done on vocabulary learning, dating back to Werner and Kaplan, 1950. Some of the methodology employed by the CVA project stems from a more recent paper, by van Daalen-Kapteins and Elshout-Mohr, 1981.

2.2.1 van Daalen-Kapteins and Elshout-Mohr (1981)

The research done by Marianne Elshout-Mohr and Maartje M. van Daalen-Kapteijns provides a strong theoretical background for the CVA project. Their work centered on analyzing the ability of different readers to determine the meaning of an unknown word from context. The results of this research shed some light on the methods by which a 'good' reader deduces this information.

An experiment they performed involved constructing several sentences with a made up word. They then had the subject read each sentence, one at a time, and describe their thoughts on what the made up word meant after each iteration.

Their findings posited that readers employ a model in the task of word learning which provides a structure for information retrieval and modification of a current hypothesis. As new information is added, readers typically apply it to their anticipated/hypothesized meaning that they have already associated with the word. Unless a significant contradiction occurs, the model remains intact as it is applied to new uses of the word.

The method we employ in SNePS bares some similarities to the system van Daalen-Kapteins and Elshout-Mohr describe. It provides a specific structure for representing particular types of information. The predictability of the locations of this information is what allows the algorithms to formulate a meaning for an unknown word. One branch of the CVA project which I won't be discussing in this paper also focuses on the revision of a known definition.

2.3 How our work relates

Previous research on CVA is made up of a very diverse body of literature including computational implementations, and psychological studies. One of the goals of our research is to unify this large disparate body of literature and incorporate elements of both disciplines into our algorithms. Ultimately the goals of the computational research will combine aspects of natural language understanding, knowledge representation and reasoning, and information extraction to produce a system that may help us better understand how we understand, while also giving us the tools to build larger, more powerful AI systems.

3 Analysis

Among the body of CVA literature, verb acquisition is a concept that has received rather little attention compared to noun acquisition. One reason for this is that nouns are simply easier to define; there is more concrete information to analyze. On the other hand, verbs represent abstract relationships. Another more obvious reason also is that there are simply more nouns than verbs in languages. Additionally, the most frequently used verbs have multiple senses;

lower frequency verbs on the other hand typically only have fewer senses. As such, in reading a body of text we are much more likely to come across an unknown noun that an unknown verb. A basic overview of the statistics of WordNet, as seen below in figure (1), exemplifies this observation.

Category	Unique Word Forms	Number of Senses
Noun	114648	141690
Verb	11306	24632

Figure 1: Statistical comparison of word forms to senses from WordNet (www.cogsci.princeton.edu/ wn/)

3.1 Analysis of contextual information

Our main goal is to determine the kinds of inferences people can make from a context, and what kind of knowledge they need to make those inferences. In order to best analyze how contextual vocabulary acquisition works, I decided to do an experiment, using several contexts containing unknown verbs, and myself as the subject. Most researchers would object to using themselves as their own subjects, as this may add an element of bias to the data generated. Also, using a larger set of subjects would provide a broader amount of data. However, in this experiment it is not the output of the process, but rather the process itself which I wished to study. Being that this is a fully cognitive process, the only subject I could gather the appropriate data from was myself.

In this experiment, I compiled a list of five verbs that I did not previously know the meaning of, or have only seen in a few obscure contexts. I then obtained several contexts for each word using the full text search available at the Project Gutenberg website (http://www.gutenberg.net). I then analyzed each context and transcribed my thoughts until I could come up with a reasonable definition for the unknown word. I then compared my results to the elements of a dictionary definition for the unknown word. At the end of the experiment I created a general set of hypotheses for how and where useful information relating to a verb is encoded in a context.

3.1.1 Contexts

In this section I will summarize the results of my experiment on a set of contexts containing ‘unknown’ verbs. In this analysis I will explicate the types of information and inferences I used to determine the meaning of an unknown word from each specific context. I will then compile the results of this analysis into an algorithm that may be merged with the current verb definition algorithm employed in the CVA project. The four verbs I have selected contexts for are augur, comport, foist, and perambulate. For each I have chosen one or two contexts which each exemplify either a different strategy that one must undertake or different information one may find to define the unknown verb.

3.1.2 Defining ‘augur’ from context

The first context I found containing the verb ‘augur’ was a very simple one from a newspaper headline:

(3.1) ‘All signs augur vigorous economy in 2004’

Even a context this simple can provide us with all the clues we need to determine the meaning of ‘augur’. In this instance all we have to define the verb is the constraints of its arguments. One major constraint is the types of actions that a ‘sign’ is capable of initiating. We can probably classify all these actions under the category of 'to make known'. At this point the context can be reevaluated as: [sign] [makes known] [abstraction] in [(future) time period]. Here we notice the constraint that the action acts on some object in a future time period. With this constraint we can presume that the action must include some component of foretelling.

Below is a slightly more complex context of 'augur'.

> (3.2) Suddenly the tempest redoubled. The poor young woman could augur nothing favorable as she listened to the threatening heavens, the changes of which were interpreted in those credulous days according to the ideas or the habits of individuals.

This context gives us several types of information. (1) causal event(s): 'the tempest redoubled'. (2) the immediate context: 'The poor young woman could augur nothing favorable'. (3) co-occurring (possibly causal) event(s): 'she listened to the threatening heavens'. (4) background information: 'the changes of which were interpreted in those credulous days according to the ideas or the habits of individuals.'

In addition to the information provided in the immediate context, this passage also gives us a selection of other actions that the agent is capable of, which may relate to the act 'augur'. First we can notice a causal chain of events: the tempest redoubles then the woman listens to the threatening heavens then the woman augurs nothing favorable. We then learn that a change in the heavens will be interpreted. Using our own background knowledge we can tell that heavens and tempest are used coreferentially in this context. Since the tempest redoubled (i.e. changed) we can tell that the heavens changed, and therefore we can conclude that the woman interprets the redoubling of the tempest, which is threatening. So now we have a more complex causal chain: the tempest redoubles then the woman listens to the tempest/heavens then the woman interprets the threatening heavens and the woman augurs nothing favorable. The important point here is that the act of 'listening' causes two events to take place. Once we break a context down this far, all we can analyze further are any parallels between the arguments of each of these phrases. How these two are related may be difficult to tell; they could be linked as cause and effect, or as two equivalent events. Either way, we can determine that 'interpret' is the closest action in the context that relates to the meaning of 'augur'.

In sum, a possible definition for 'augur' that we can infer from just this context is that it is closely related to 'interpret'; interpreting something threatening results in something bad/not favorable. If we combine this with the results of the analysis on the previous context of 'augur' we can include 'foretelling' as a related action, and a 'sign' as a possible effector.

According dictionary.com, which contains definitions from a large number of dictionaries, 'augur' has the following definition:

> To predict, especially from signs or omens.

We can see here that most of the inferences that I made during this experiment, when I did not know the meaning of the word 'augur' produced all the necessary components of its meaning. Later in this paper I describe a computational implementation of the passage in (3.2), and incorporate some of the inference rules I discussed above. From this implementation we will be able to see what aspects of the verb algorithm need to be augmented to return a well formed definition.

3.1.3 Defining 'comport' from context

Example (3.3) contains a context of the word 'comport'.

> (3.3) There were, also, divers ladies in New York, Newport, and elsewhere, and celebrated for their palatial homes, their jewels, and their daughters, who were anxious to know how Bellew would comport himself under his disappointment. Some leaned to the idea that he would immediately blow his brains out; others opined that he would promptly set off on another of his exploring expeditions, and get himself torn to pieces by lions and tigers, or devoured by alligators; while others again feared greatly that, in a fit of pique, he would marry some 'young person' unknown, and therefore, of course, utterly unworthy.

The most significant element of this context lies in the consequences of the word 'how' in '...who were anxious to know how Bellew would comport himself under his disappointment.' The sentence following then provides us with three possible methods for how someone who is disappointed would engage in comporting himself. These events include committing suicide, going and being killed on an expedition, or getting married to someone unworthy. Analyzing the semantics of these phrases, it is probably impossible to generalize and develop a single action that they all fall under, except to say that they are all some sort of act that the agent of 'comport' may engage in sometime in the future.

In sum, a definition for 'comport' that we can form from just this context might be: 'performing some act as a result of an emotional state'.

According dictionary.com, 'comport' has the following definition:

> To conduct or behave (oneself) in a particular manner

From this we can see that some basic components of the meaning of 'comport' were indeed available in the context. The components of 'act' and 'emotional state' fit fairly closely to the components of 'behave' and 'manner' in the actual definition.

3.1.4 Defining 'foist' from context

Example (3.4) contains a context of the word 'foist'.

> (3.4) I was to look for a watch that I knew he hadn't lost, and was to receive 200 pounds if I found it. It seemed to him to be a very good bargain, as, indeed, it was, from his point of view, feeling, as he did, that there never having been any such watch, it could not be recovered, and little suspecting that two could play at his little game of deception, and that under any circumstances I could foist a ten-shilling watch upon him for two hundred pounds. This business concluded, he started to go.

In this context we are mainly provided with background information on the situation before the action 'foist' takes place. We can gather a central theme from this information, which seems to be one of 'deception'. We might be able to say that this is possibly a component of the meaning of 'foist'. The immediate context is what we can form the most inferences from: 'I could foist a ten-shilling watch upon him for two hundred pounds.' Judging from the arguments 'foist' subcategorizes for the some of the same argument types as verbs like 'sell' (i.e. an item, 'for', and a sum of money); however it also subcategorizes for a prepositional phrase headed by 'upon'. The problem with prepositions like 'upon' is that they can be used in many different ways, the meaning of which is dependent on the meaning of the verb that it is linked to. The most general description of 'upon' is that it emphasizes directionality toward the undergoer.

In sum, a possible definition for 'foist' that we can infer from just this context is that it may involve deception and it incorporates the act of selling and performing some imposition upon the undergoer.

A second context for 'foist' is given in example (3.5)

(3.5) They had their little argument. I lit a cigarette and let them argue. In such cases, every married couple has its own queer and private and particular and idiosyncratic way of coming to an agreement. The third party who tries to foist on it his own suggestion of a way is an imbecile.

This passage gives us the word ‘foist’ in a very different context. It subcategorizes here for a different set of arguments, and the superordinate categories of the arguments are also different than in the previous context. A logical conclusion from this observation may be that ‘foist’ has an additional meaning that is being used here. The basic structure present here is roughly [an agent] can foist [a concept] on [something]. One thing that is still in this context is the incorporated meaning that ‘on’ adds to the relationship of the arguments. As a result, another action that might fit into the syntax and semantics of this passage is ‘force’.

According dictionary.com, ‘foist’ has the following definitions:

a. To pass off as genuine, valuable, or worthy
b. To impose (something or someone unwanted) upon another by coercion or trickery.
c. To force onto another.

We can see here that most of the inferences made for each context were indeed correct; (a) and (b) include features that were inferred from the analysis of (3.4), and (c), which is somewhat different, paraphrases what was inferred from (3.5).

3.1.5 Defining ‘perambulate’ from context

Example (3.6) contains a context of the word ‘perambulate’.

(3.6) Katharine would calculate that she had never known her write for more than ten minutes at a time. Ideas came to her chiefly when she was in motion. She liked to perambulate the room with a duster in her hand, with which she stopped to polish the backs of already lustrous books, musing and romancing as she did so. Suddenly the right phrase or the penetrating point of view would suggest itself, and she would drop her duster and write ecstatically for a few breathless moments; and then the mood would pass away, and the duster would be sought for, and the old books polished again.

This context gives us several types of information. (1) Background information on the agent: ‘Ideas came to her chiefly when she was in motion’. (2) The immediate context: ‘She liked to perambulate the room’. (3) Causal information: ‘She liked to perambulate the room... ...Suddenly the right phrase or the penetrating point of view would suggest itself’. The most important information lies in the causal chain of events presented. Summarized, they are: Ideas came to her when she was in motion, when she perambulated the room, ideas came to here. We can conclude that perambulating involves the agent being in motion. From our own background knowledge we can presume that since she is in a room she is moving by walking.

According dictionary.com, ‘perambulate’ has the following definitions:

To walk through.
To walk about; roam or stroll.

From this we can see that the inference of ‘walking’ was indeed correct. We can also look back and see what other inferences could have been made. the related actions, ‘roam’ and ‘stroll’ typically include components of movement without a particular destination. This may be inferred from the passage in (3.6) from the general mood of the agent and possibly from the theme of the entire passage. Retrieving this information by pulling out key descriptive words and phrases, and generating further inferences from them. An example of this may be: ‘duster in her hand, she stopped to polish the backs of already lustrous books, musing and romancing as she did so’. In retrospect these are all important clues for the manner in which the act was performed.

3.2 Observations

From these findings I determined that there are two main types of contextual information immediately available in a context; the immediate context of the verb, which includes its arguments and the verb phrase’s adjuncts, and the wider context in which is generally present a causal chain of events incorporating the unknown action, or a gradual development of description of a state which the unknown action falls into.

The immediate context provides us with the possible constraints of the verb’s arguments. The syntactic information present here tells us what types of arguments the verb subcategorizes for. The semantics of the arguments give us the potential class inclusions for the agent and undergoers. We can also compile an entire host of attributes from which we may infer various constraints on the verb, such as properties of the agent, other actions the agent can perform, and other actions that can be performed on the object. The immediate context will also contain any additional modifiers of the action, such as adverbs or adverbial adjuncts that add information to the event occurring.

Another piece of information that may occur in the immediate context is the presence of a conjoined event. From the contexts I analyzed, I found that there are two possible interpretations of conjoined events. Consider the following examples.

a. It rained in Dallas and snowed in New York.
b. I put on my shoes and went outside.

In (a) we would interpret the two events as occurring simultaneously, whereas in (b) we can identify that the two events occurred in sequence. In each of these examples, the information in one event provides some information that is part of the other.

In (a) we can tell that the two events share some components of meaning, namely ‘precipitation’. Additionally, the syntax of the two clauses is identical, as is the class membership of the arguments. By constrast, in (b) we can tell that the first clause is the antecedent and the second clause the consequent. Additionally, the class membership of the verbs’ arguments differs significantly, one being ‘shoes’, the other being ‘outside’.

Identifying these constructions in a representation may be useful for identifying the relationship of an unknown verb to other actions in its context. Since most dictionaries define verbs in terms of other related actions, this is probably the most important piece of information that can be found from a context. The wider context of the unknown verb may represent either the continuous development of a description of a state, or a causal chain of events which the unknown action falls into. It is difficult to develop a general set of rules that relates this information to the unknown action because the structure and content can vary greatly. An analysis of this component may yield different observations from context to context. Sometimes the connection that information in the wider context has to the unknown event is extremely subtle, and only accessible if the reader does know the meaning of the ‘unknown’ word.

Sometimes the wider context will define either background knowledge, or domain-specific rules of inference that apply only to the information in the context itself. These are often the types of things that a generic algorithm cannot account for; instead it may be up to the parser (or the person representing the passage in SNePS) to identify this information.

3.3 More Analysis

The goal we are striving for is to enhance the verb algorithm so that it returns a potential definition of an unknown verb. To do this, we must first analyze the manner in which dictionaries define verbs.

Example (3.7) lists definitions (from www.dictionary.com) for the words analyzed in section 3.1 as well as some of the other words that have been studied in contexts for the CVA project. I will use this data to analyze the types of information that form a suitable definition of a verb.

(3.7)

Augur: To predict, especially from signs or omens.

Comport: To conduct or behave (oneself) in a particular manner

Foist:

- **a.** To pass off as genuine, valuable, or worthy
- **b.** To impose (something or someone unwanted) upon another by coercion or trickery.
- **c.** To force onto another

Perambulate:

- **a.** To walk through.
- **b.** To walk about; roam or stroll.

cripple: To disable, damage, or impair the functioning

fold: (itr.)

- **a.** To close, especially for lack of financial success; fail.
- **b.** To give in; buckle.

proliferate:

- **a.** To grow or multiply by rapidly producing new tissue, parts, cells, or offspring.
- **b.** To increase or spread at a rapid rate.

pry:

- **a.** To raise, move, or force open with a lever.
- **b.** To obtain with effort or difficulty.

quail:

- **a.** To shrink back in fear; cower.
- **b.** To cause to fail in spirit or power; to quell; to crush; to subdue.

From this data, I propose the following as a list of the information present in a dictionary definition in rough order of importance:

Argument structure (transitivity)
Similar actions
Similar action plus spatial relation(s) (e.g. prepositional phrase)
Manner of the performance of the action
Causes

Effects
Type of Object(s)
Type of Instrument(s)
Type of Agent

Transitivity is listed as the most important because this is what typically separates one sense of the verb from another. Similar actions, and similar actions combined with prepositional phrases are the most common way of specifying the sense of the verb. Manner of performance is next, which is important for describing properties of the action. Information on cause and effect provide useful information on the change of state created by the verb, however it usually leaves us without any actual representation of the process itself. An ideal definition would contain cause, effect, and manner in order to fully describe the meaning of a verb. Class membership of the verb's arguments is typically described in terms of the superordinate category that possible arguments will fall under, and as such does not provide very detailed information all that often.

One general observation from this is that most definitions ignore the agent, most likely leaving it as an unspecified, generic 'causal agent'. However one thing to take into consideration is that all verbs in existence that require an agent are able to be initiated by a human agent. As such, information relating to the agent may not be as important as I previously suspected. The only word I found that included a possible agent type in its definition was 'hibernate'.

In conclusion of this analysis, dictionary definitions are actually quite simplistic in their content and construction. The problem is that the most important pieces of information that make up a good definition are the hardest to pull out of a context; namely, related actions and manner. We are usually pretty good at identifying the categories of the verb's arguments. In fact, this represents the bulk of the information that the current verb definition algorithm extracts. As such, inferring possible related actions from this information is an important step for this algorithm to include, especially when no candidates for similar actions appear in a context.

4 Implementation

4.1 Implementation of the 'augur' passage.

In order to test the enhanced verb algorithm I implemented a representation of a context of 'augur'. The context, which appeared above as example (3.2) is repeated here as (4.1). Developing a representation of this passage allowed me to do several things: (1) identify the best way to represent each type of information analyzed in the previous section, (2) determine the required background knowledge and rules needed to understand the passage, and (3) to compare my own initial hypothesis on the meaning of 'augur' to the output of the algorithm.

> (4.1) Suddenly the tempest redoubled. The poor young woman could augur nothing favorable as she listened to the threatening heavens, the changes of which were interpreted in those credulous days according to the ideas or the habits of individuals.

One thing to note is that all the decision-making is done by the inference rules defined in SNePS. The only purpose of the verb algorithm is to search through a set of key locations in the network and summarize the information it finds in a manner that best represents a definition of the unknown word. At this stage, the verb algorithm is able to pull most of the relevant information out of the network, and provide some basic generalizations of contexts the word might be used in. It currently has not reached the level of functionality needed to output a potential dictionary definition for the word.

The SNePS representation for the passage given in (4.1) includes case frames (i.e. proposition node + argument arcs) with the syntax and semantics specified in figures (2) and (3). The importance of specifying the semantics of case frames used lies in the fact that the arc labels are completely meaningless to the system. The only relevance they have is in how we decide to use them in conjunction with other inference rules. The labels are merely a mnemonic for us to remember their function by.

4.1.1 Knowledge Representation for the passage

The representation for the passage given in (4.1) includes the following propositions, paraphrased in English.

General Background Rules:

If action A is performed by agent Y on object Z, and action B is performed by agent Y on object Z, then A and B are similar.
If action A and action B are similar, and A causes C, then B causes C as well.
If X and Y are equivalent, and X has property A, then Y has property A.
If X and Y are equivalent, and X does act A, then Y does act A.
If X and Y are equivalent, and X is a member of class A, then Y is a member of class A.
If X and Y are equivalent, and X is a subclass of A, then Y is a subclass of A.

Background Knowledge for the passage:

Tempest and heavens are equivalent.
The woman is a member of the class human.
'Nothing favorable' has the property of being not favorable.
The days are credulous.
The class hierarchy: tempest ⇒ windstorm ⇒ storm, violent storm ⇒ atmospheric phenomenon ⇒ physical phenomenon ⇒ natural phenomenon ⇒ phenomenon.
The class hierarchy: human ⇒ animate thing ⇒ physical object ⇒ entity.
The class hierarchy: days ⇒ time period ⇒ fundamental quantity ⇒ amount ⇒ abstraction.
The class hierarchy: quality ⇒ attribute ⇒ abstraction.
The class hierarchy: interpret ⇒ explicate ⇒ inform ⇒ communicate ⇒ interact ⇒ act.
The class hierarchy: listen ⇒ perceive — concentrate ⇒ think.
The class hierarchy: redouble ⇒ escalate ⇒ increase ⇒ change.

Information from the passage:

The tempest redoubles.
The heavens are threatening.
The woman is young.
The woman is poor.
All people interpret an instance of change in the heavens.
An instance of change is a member of the class change.
If the days are credulous, then all people are credulous.
If a person is credulous, they are 'disposed to believe on little evidence'.

<table>
<tr><th>SNePS case frame</th><th>Semantics</th></tr>
<tr><td>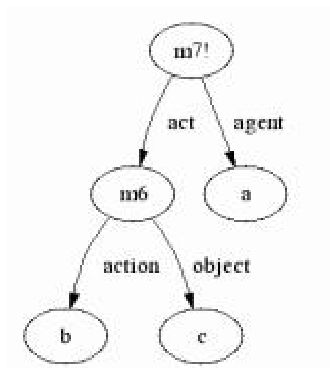
</td><td>(m7!) is the proposition such that [[a]] engages in action [[b]] with respect to object [[c]]. Basically this case frame represents the structure of a transitive verb.</td></tr>
<tr><td>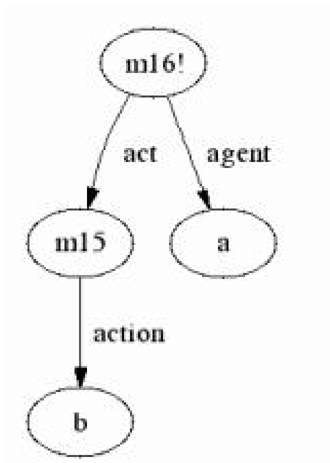
</td><td>(m16!) is the proposition such that [[a]] engages in action [[b]]. Basically this case frame represents the structure of an intransitive verb.</td></tr>
<tr><td>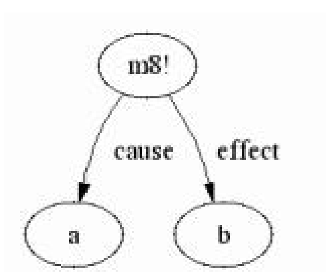
</td><td>(m8!) is the proposition such that [[a]] is the cause of [[b]].</td></tr>
<tr><td>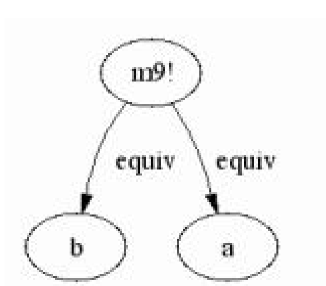
</td><td>(m9!) is the proposition such that [[a]] is equivalent to [[b]].</td></tr>
</table>

Figure 2: Case frames used in the SNePS representation of the ‘augur’ passage.

<table>
<tr><th>SNePS case frame</th><th>Semantics</th></tr>
<tr><td>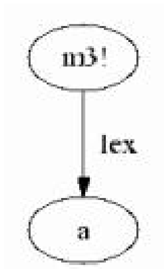
</td><td>(m13!) is the concept expressed in English as ‘a’.</td></tr>
<tr><td>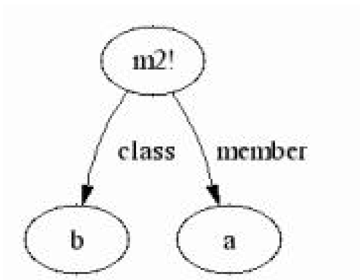
</td><td>(m9!) is the proposition such that [[a]] is a member of the class [[b]].</td></tr>
<tr><td>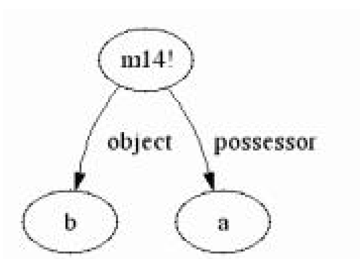
</td><td>(m14!) is the proposition such that [[a]] possesses [[b]].</td></tr>
<tr><td>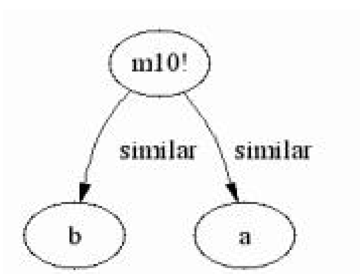
</td><td>(m10!) is the proposition such that [[a]] is similar to [[b]].</td></tr>
<tr><td>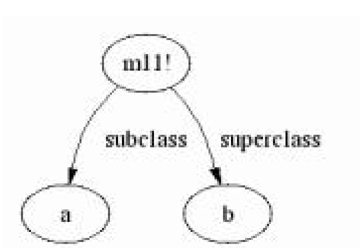
</td><td>(m11!) is the proposition such that [[a]] is a subclass of [[b]].</td></tr>
</table>

Figure 3: Case frames used in the SNePS representation of the ‘augur’ passage.

If something redoubles, then it possesses an instance of change.
The proposition that the woman listened to the change in the heavens caused the Proposition that the woman augurs nothing favorable.

Additional inferences from the passage:

If a person interprets an instance of change and the days are credulous, then the instance of change will be interpreted to be equivalent to a bad omen.
If something is equivalent to an instance of change, and the act of interpreting an instance of change caused some event, then that thing is also the cause of the event.
A bad omen is not favorable.

4.2 Results of the 'augur' representation

The output of the verb algorithm when running the augur demo file is displayed below.

```
* * * Defining the verb augur * * *


Arguments of the verb:
  (agent object)

Transitivity:   transitive


* * * Basic Findings * * *


  possible cause of augur is:
      bad omen,
      heavens,
      instance of change,
      tempest,

  possible actions performed by agent of augur is:
      augur,
      interpret,
      listen,

  possible actions performed on object of augur is:
      augur,

  possible action that is the cause of augur is:
      listen,
      interpret,
```

```
 possible property of verb is:
     unknown,



 possible membership of agent is:
     human,

 possible superordinate of agent is:
     animate thing,

 possible property of agent is:
     disposed to believe on little evidence,
     credulous,
     poor,
     young,


 possible superclass of object is:
     quality,

 possible property of object is:
     not favorable,

* * * generalizations from this context * * *


 A {human} can augur
        Something that is a subclass of {quality}

 A {human} can augur
        Something with the properties {not favorable}

 A {animate thing} can augur
        Something that is a subclass of {quality}

 A {animate thing} can augur
        Something with the properties {not favorable}

 Something with the properties {disposed to believe on little evidence, credulou
        Something that is a subclass of {quality}

 Something with the properties {disposed to believe on little evidence, credulou
```

```
          Something with the properties {not favorable}
nil

 CPU time : 0.07

*
```

The output of the verb definition algorithm is divided into three parts: first there is syntactic information derived from the context. This currently includes some basic information on the types of arguments the verb takes as well as its transitivity.

The second part of the output lists all the relevant information of the arguments in the context. From this part of the output we can assess the accuracy of the algorithm and the inference rules employed in the representation. Overall, the verbs ‘listen’ and ‘interpret’ are given as the closest actions to ‘augur’. Another important piece of information is the possible causes of ‘augur’: bad omen, heavens, instance of change, tempest. ‘bad omen’ and ‘instance of change’ are fairly good guesses for the cause; the inclusion of ‘heavens’ and ‘tempest’ is a conclusion specific to this context. Under properties of the agent, it includes ‘credulous’, which may add to the ability for someone to infer the meaning of ‘augur’.

The third part of the output takes the information that was found and constructs a possible set of new contexts that the verb may occur in. Probably the most informative generalization we can get out of this particular output is that *something that is disposed to believe on little evidence can augur something that is not favorable*. One immediate improvement could be to include causes or effects in these generalizations. With that we may be able to add that *something that is disposed to believe on little evidence can augur something that is not favorable from a bad omen*, which is somewhat closer to the actual meaning.

Obviously these generalizations are only a first step to outputting a well formed definition of the word. Improving this aspect of the verb definition algorithm could likely be an object of research for students working on the CVA project in the future.

5 Conclusions

From the analyses and implementation of the algorithm presented in this paper I will propose a possible word learning strategy for acquiring verb meanings from context.

1. Once an unknown verb is encountered, determine the arguments that it takes.

   (a) Determine the basic class membership of the object(s) and any instruments.

   (b) Identify the prepositions that head the verb’s arguments.

   (c) Using both of these pieces of information, hypothesize from background knowledge of language what other verbs may fit into the syntactic and semantic context. (this is probably the most common strategy employed to some extent by most people, however many people will tend to ignore either the syntax or semantics and make a guess that fits either one or the other)

2. To further explicate the process of (1c), identify argument constraints on the verb:

   (a) Categorize all actions that the agent is capable of initiating. Return the lowest level superordinate category that includes all (or most) of those actions.

- **(b)** For transitive verbs, categorize all the actions that can be performed on the object(s). Return the lowest level superordinate category that includes all (or most) of those actions.
- **(c)** Return the intersection of the lists of all actions that fall under the categories returned by (2a) and (2b).

3. Identifying cause and/or effect of the unknown verb:

- **(a)** Determine what event is the immediate antecedent of the event containing the unknown verb.
- **(b)** Determine if there is a causal chain of events that the unknown action belongs to (i.e. a sequence of more than 2 events which include the unknown action).
- **(c)** Determine any co-occurring events with the unknown action.
- **(d)** Infer any causal or equivalent events.
- **(e)** Determine any resulting events from the unknown action.
- **(f)** Determine if the verb is part of a resultative construction.
  - **(i)** possible clues might be class membership of the final argument.

4. Identify the manner in which the action is performed:

- **(a)** Return the superordinate categories of the method(s) in which the action is performed.

5. Thematic overtones.

- **(a)** Identify the overall topic or tone of the passage. (This may include hints on the manner of the action. Sometimes it is based on the overall goals of the participants in the action from other portions of the text)

6. Compile all of the information gathered.

- **(a)** Identify any other known actions that contain those components.

Additional research will need to be done in order to validate or refute the effectiveness of this strategy. The immediate next step would be to implement parts of it in the verb definition algorithm, and implement representations of additional contexts in light of the analyses described in this paper.

The above steps may also be tested as a vocabulary learning strategy to gather feedback on its effectiveness for improving human vocabulary acquisition. The results of this research can then go back into development of the computational portion of the project.

References

[1] Carbonell, Jaime G. 1979. *Towards a Self-Extending Parser*, Proceedings of the 17th Annual Meeting of the Association for Computational Linguistics (University of California at San Diego). (Morristown, NJ: Association for Computational Linguistics): 3-7.

[2] Dictionary.com. 2004. *http://www.dictionary.com* Lexico Publishing Group, LLC.

[3] Hastings, Peter M., and Lytinen, Steven L. 1994. *Objects, Actions, Nouns, and Verbs*, Proceedings of the 16th Annual Conference of the Cognitive Science Society. (Hillsdale, NJ: Lawrence Erlbaum Associates): 397-402.

[4] Granger, Richard H. 1977. *Foul-Up: a Program that Figures Out Meanings of Words from Context*, Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI-77, MIT) (Los Altos, CA: William Kaufmann): 67-68.

[5] Project Gutenberg. 2004. *http://www.gutenberg.net/*

[6] Rapaport, William J., and Ehrlich, Karen. 2000. *A Computational Theory of Vocabulary Acquisition,* in Lucja M. Iwanska and Stuart C. Shapiro (eds.), Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language (Menlo Park, CA/Cambridge, MA: AAAI Press/MIT Press): 347-375.

[7] Rapaport, William J., and Kibby, Michael W. 2002. *ROLE: Contextual Vocabulary Acquisition: From Algorithm to Curriculum,*

[8] Shapiro, Stuart C., and Rapaport, William J. 1987. *SNePS Considered as a Fully Intensional Propositional Semantic Network,* in Nick Cercone & Gordon McCalla (eds.), The Knowledge Frontier: Essays in the Representation of Knowledge (New York: Springer-Verlag): 262-315.

[9] Schank, Roger C., and Rieger, Charles J., III. 1974. *Inference and the Computer Understanding of Natural Language,* Artificial Intelligence 5: 373-412.

[10] van Daalen-Kapteijns, M.M., and Elshout-Mohr, M. 1981. *The Acquisition of Word Meanings as a Cognitive Learning Process,* Journal of Verbal Learning and Verbal Behavior 20: 386-399.

[11] Werner, Heinz, and Kaplan, Edith. 1950. *Development of Word Meaning through Verbal Context: An Experimental Study,* Journal of Psychology 29: 251-257.

[12] Wiemer-Hastings, Peter; Graesser, Arthur C. and Wiemer-Hastings, Katja. 1998. *Inferring the Meaning of Verbs from Context* in Proceedings of the 20th Annual Conference of the Cognitive Science Society.

[13] WordNet. 2004. *http://www.cogsci.princeton.edu/ wn/* Princeton University.