The Acquisition of Common Sense Knowledge by Being Told: An Application of NLP to Itself

Fernando Gomez

School of EECS
University of Central Florida, Orlando, FL 32816
Harris Center
gomez@eecs.ucf.edu

Abstract. This paper shows how the knowledge of a semantic interpreter can be bootstrapped for other semantic interpretation tasks. Methods are described for automatically acquiring common sense knowledge and for applying this knowledge to noun sense disambiguation. Ordinary concepts are described by several plain English sentences that are parsed and semantically interpreted. The semantic interpreted sentences are stored under these concepts to be used for semantic interpretation tasks. This paper explains the description of the concepts, the interpretation of the sentences and two algorithms for noun sense disambiguation that use the acquired knowledge.

1 Introduction

In [2], a method for defining verb predicates and an implemented algorithm [4] that resolves the verb predicate, its semantic roles and adjuncts is explained. The approach has proceeded first by defining verb predicates for WordNet verb classes [1], and then by defining verb predicates for individual verbs with a high degree of polysemy. The semantic roles of the predicates are linked to the selectional restrictions (categories in WordNet ontology for nouns) and to the grammatical relations that realize them. The selectional restrictions of the predicates are grounded on the WordNet ontology for nouns [10], whose upper level ontology has been modified and rearranged [3] by using the feedback obtained by testing the predicate definitions. The resulting work is immense: over 3000 verb predicates have been built. A corpus of over 3000 semantically interpreted sentences has been automatically created with the semantic interpreter. We have used the semantic interpreter in order to extend it in several important ways. There are several semantic tasks that the interpreter does not solve or solves them partially. One of them is noun sense disambiguation (NSD). In many cases, the selectional restrictions of the verb predicates cannot resolve the noun senses of the semantic roles, although it may narrow them down to a small set of possible senses. For instance, in a sentence such as “The batter dropped the bat” the selectional restrictions for the predicate of *drop-something-physical* select the first sense of “batter.” They also rule out the third sense of “bat,” (a turn batting in baseball), which has *activity* as its hypernym. But, the semantic interpreter

E. Kapetanios, V. Sugumaran, M. Spiliopoulou (Eds.): NLDB 2008, LNCS 5039, pp. 40–51, 2008.

cannot decide between the other five senses of "bat," which are all physical things, because *physical-thing* is the selectional restriction for the *theme* of *drop-something-physical*. The output of the semantic interpreter for that sentence is depicted in Figure 1. The interpreter groups the noun senses by ontological category: bat1 is baseball bat, and bat5 is a cricket bat, which have *equipment* as their upper-level ontological category. The noun senses are WordNet senses. Concepts without number senses correspond to Gomez's WordNet upper-level ontology.

```
((SUBJ ((DFART THE) (NOUN BATTER)) ((BALLPLAYER1 BATTER1)) (AGENT))
 (VERB DROPPED ((MAIN-VERB DROP DROPPED)) DROP-SOMETHING-PHYSICAL
      (DROP1 DROP2) SUPPORTED BY 2 SRS)
 (OBJ ((DFART THE) (NOUN BAT))  ((EQUIPMENT BAT1 BAT5) (PLACENTAL1 BAT2)
      (INSTRUMENTALITY BAT4 BAT6))  (THEME)))
```

Fig. 1. Output for "The batter dropped the bat"

In the sentences, "Farmers like plants," "The doctor removed the appendix in an operation," "Several demonstrators were injured in the demonstration," none of the senses of "plant," "doctor" "appendix," and "operation" can be ruled out on the basis of selectional restrictions, or by similarity in the WordNet taxonomy. However, most humans will tell us that, in those sentences, "plant" is a life form, and not a industrial plant, or an actor, etc., and that "appendix" is not the appendix of a book, but an animal body part, that "doctor" refers to a physician and not to a theologian, that "operation" is a surgical operation, and not a military operation or a business operation, etc. Humans determine these senses using commonsense knowledge which is as basic as that used in the selectional restrictions of verb predicates. The aforementioned examples do not provide much context, but sufficient to establish the senses of the nouns. The aim of this research is to let the semantic interpreter acquire this knowledge by being told and use it for NSD, prepositional attachment, discourse, etc.

The glosses in dictionaries have been used for noun sense disambiguation [6]. There have been also efforts to produce some kind of logic form transformation of WordNet glosses [11]. However, many of the glosses in WordNet and in dictionaries do not lend themselves to having knowledge extracted from them by a program, except in a superficial manner, because they require much knowledge to be understood. We understand many of the dictionary glosses because we already know a lot about the concepts expressed in them. A serious problem with the glosses is that lexicographers approach them as succinct definitions of concepts intended to capture their essential aspects. As a result of this, in many instances the language of glosses contains many intensional terms [13] that do not indicate how the words are used [15] in ordinary language. For instance, the WordNet gloss for cell2 is "the basic structural and functional unit of all organisms." The meaning of the terms "basic structural unit" and "basic functional unit" do not lend themselves to be easily acquired by a program because of their abstract content: two abstract adjectives predicated of an abstract noun, "unit." It is unlikely that a child will be introduced to the concept of cell2 that way. However, a sentence such as "All living things are made of cells" provides a better introduction to the concept of cell2 because the concept "living thing," as well as that of cell2 have well established denotations to physical objects. Hence, it becomes easy to understand sentences such as "Trees are made of cells," or compound nouns such as "plant cells" by subsumption between the concepts in the new sentences and the concepts in the defining sentence, or gloss. Of course, we will need other defining sentences, or glosses, to provide the basic aspects of the concept of cell2. Likewise, the WordNet gloss for demonstration3, "a public display of group feelings," is very good if one already knows the meaning of demonstration3. But, the gloss is not very helpful if one is acquiring the concept for the first time. Instead of that definition, we prefer glosses that indicate the events normally associated with demonstration3 e.g., "In a demonstration, humans gather in streets or plazas to protest," "Sometimes in a demonstration, people may be hurt, or physical objects damaged," "A riot may commence during a demonstration" etc. In summary, stay away from intensional terms and do not try to convey many aspects of a concept in one single sentence and assume that you are conveying these concepts to someone who does not know anything about them except their ontological classification as provided by WordNet. In this paper, we show how to provide glosses for WordNet word senses, how to parse and semantically interpret them and, then, use them for noun sense disambiguation. In the next section, we explain the definition of the glosses and their semantic interpretation. In sections 3, we deal with the relation of verb selectional restrictions and noun senses. In sections 4, and 5, we explain the algorithms for noun sense disambiguation. Sections 6 provides the testing of the algorithms, and sections 7 and 8 explain related research and conclusions, respectively.

2 Acquiring Common Sense Knowledge about Ordinary Concepts

There is no difference between a noun sense and a concept. Ambiguous words stand for various concepts, or noun senses. Thus, learning about a new sense is not different from learning about concepts. What we are describing in this paper is the acquisition of basic knowledge about concepts, and its application to an aspect of semantic interpretation, namely noun sense disambiguation. But, this knowledge can be used for other aspects of semantic interpretation, or for other applications. We have used the word "glosses" to refer to the sentences describing some of the main aspects of a concept, or noun sense. The acquisition is as simple as typing some sentences. Suppose that one wants the system to acquire some basic knowledge about pot4, "a container in which plants are cultivated." One types (acquire pot4) to lock the senses of "pot" to pot4, and one starts typing some sentences, one at a time. For instance, one may type: "People cultivate plants in a pot," "Soil is put in pots," "People plant plant parts and plants in a pot," etc. In those cases in which the system is unable to determine the senses of some of the nouns in the gloss (henceforth, GL), the user may tell the system the correct senses by typing (*refer-to* $sense_1, sense_2...sense_i$). One line fixes all senses at once. But in most cases, the system is able to determine the senses of the nouns in the glosses. For instance, out of the three glosses for pot4, the system only fails to determine the sense of "soil." For some concepts, three or four sentences are sufficient for providing the main aspects of that concept, while other concepts may need more sentences. The system parses and interprets most sentences in one or two seconds. Next, one provides glosses for the other senses of "pot." Glosses for concepts which do not denote physical objects are provided in the same way. For instance these are some of the glosses for operation7 (a surgical operation). "Doctors performed operations," "In an operation, doctors operate on humans or animals," "In an operation, some body parts are removed, or replaced," "Some diseases may be only cured by means of an operation," "Some patients may die as a result of an operation," "Most operations take place in hospitals or clinics." This set of glosses provides a frame-like or script-like type of knowledge.

3 Verb Selectional Restrictions and Noun Senses

If one enters the sentence "She ate the dates with a fork," our interpreter will select (edible-fruit1 date8), and (cutlery2 fork1). The verb predicate, the noun senses and the semantic role are all solved. But, there are many sentences for which the semantic interpreter cannot solve the noun senses based on the selectional restrictions for the semantic roles. In a framework in which the semantic interpreter is activated first to determine verb meaning (verb predicate) and semantic roles, the following cases caused by noun polysemy may occur:

- (a) The semantic interpreter resolves the verb predicate and the semantic roles, but it cannot decide on some noun senses.
- (b) The semantic interpreter is unable to narrow the verb predicates to one.
- (c) The interpreter comes up with more than one semantic role for the same grammatical relation.

The sentence, "He put the batter in the refrigerator," is a good example of case *(a)*. The verb predicate and semantic roles are solved, but the semantic interpreter cannot decide between the two senses of "batter," a baseball player or a flour mixture. That would not be the case for the sentence, "The batter cleaned/fixed the refrigerator."

```
[to-loc(xor natural_elevation1 peak5 physical-thing) (obj)
         (conveyance3) ((prep aboard))
        (xor natural_elevation1 peak5  physical-thing)((prep up to))
       (physical-thing)((prep onto into))]
```

Fig. 2. A definition of the role *to-loc* for one of the predicates of "climb"

An example of case *(b)* is the sentence, "The runners are raised in a nursery for one growing season." The interpreter comes up with the following verb predicates for "raise" in order of preference: RAISE-FARM (farm plants and/or animals), BRING-UP (to educate somebody), RAISE-SOMETHING (to lift something). The semantic roles for these verb predicates are all the same, namely, *theme, at-loc,* and *duration.* The noun senses selected for the *theme* of RAISE-FARM are runner8 (a fish) and runner5 (a stolon, a plant part). The noun senses selected for the *theme* of BRING-UP are runner1, runner2, runner3, runner4 runner6, all of which have *person* as their hypernym. The noun senses selected for the *theme* of RAISE-SOMETHING are all the 8 senses of "runner." The noun senses selected for the roles *at-loc* and *duration* are the same for all the verb predicates. The sense for "nursery" (a room for a baby, and a place to cultivate plants) cannot be resolved by the semantic interpreter in any of the predicates. "Growing season" is not ambiguous.

For an example of case *(c)* consider the sentence "A swell lurched the catamaran towards the reef." For this sentence, the semantic interpreter resolves the verb predicate for "lurch," but it cannot decide between the roles, *agent* and *inanimate-cause* (a causal agent other than a human, a social group, or an animal). For the *agent*, the interpreter selects the sense swell4 (a dandy, a person) and for the *inanimate-cause* the senses swell1 (a wave) and swell2 (a natural elevation1).

One of the difficulties in defining selectional restrictions for semantic roles is staying away from over-generalization and over-specification of the ontological categories in the selectional restrictions. Over-specification results in failing to identify some semantic roles, or over-narrowing the noun senses; while over-generalization results in not selecting between different noun senses and/or not distinguishing between verb senses, or verb predicates. Figure 2 depicts the definition of the role *to-loc* for one of the predicates of the verb "climb" when it means traveling upwards. The syntax for the roles is:

(role (slr) (grs) (slr) (grs) ... (slr) (grs))

Where *slr* stands for any number of selectional restrictions, and *grs* for any number of grammatical relations. When the semantic role is realized by prepositions, the prepositions are put in a list preceded by the word "prep." If there is more than one selectional restriction, the list must be preceded by "xor" or "xand". These entries indicate how the items in the list should be matched. In Figure 2, the first sublist contains the selectional restrictions (WordNet ontological categories) for the grammatical relation, direct object. Selectional restrictions are always matched from left to right. The entry "xor" in the sublist means that as soon as an ontological category in the sublist is matched to the head noun of the grammatical relation, the others ontological categories are not tried. Thus, for the sentence, "She climbed the hill," the system will select hill1, a natural elevation, for the *theme*. The entry "xand"in the sublist means that all ontological categories in the sublist are matched, but the senses are preferred in the order in which they are matched. Thus, if the role *to-loc* were defined using "xand," the senses selected for "hill" in "She climbed a hill" would be hill1, hill2 (mound4) and hill3 (mound1).

The definitions for the semantic roles aim at striking a balance between over-specification and over-generalization. In constructing the selectional restrictions for a semantic role using an "xor" entry, one needs to be very careful with those words whose senses range over several ontological categories in the list. Consider the sentence, "She climbed the table," and the *to-loc* role defined with an "xor." The system will select the sense table4 ("flat tableland with steep edges") which is natural elevation in WN, and will exclude the other senses of "table," which may not be correct. Suppose that one writes the following ontological categories for the *instrument* of "kill" when it means *cause-to-die*:

(instrument(xor weaponry1 external-body-part1 physical-thing) ((prep with)))

Assume that the sentence to be interpreted is "Samson killed the lion with the arms." That definition will select the sense of arm3 (a weapon) for "arm" and exclude its other senses.

Now, consider the sentence, "The dish is made of chicken." Humans have no trouble in determining that "dish" is not dish1 (a piece of dishware) or dish5 (an antenna etc.), but dish2 (an item of prepared food). The distinction between dish1 and dish2 is a difficult one because we put food in dish1 or dish2. Similarity in the WN taxonomy does not help. Solving this problem by using selectional restrictions will require a very specific predicate, *made-of-food*. However, we can tell the system that "Dishes (dish2) are made of food," then parse and interpret the sentence, and use the interpreted sentence for deciding between the senses of "dish." The interpreted sentence can be inherited by all hyponyms of dish2. Consider the sentence "The bowl contains punch." No selectional restrictions or similarity in the WN taxonomy will help in this sentence either. However, the simple gloss "Dishes (dish1) contain food," will solve the senses of "bowl" and "punch" because glosses are inherited by all hyponyms of the synset where they are stored. Consequently, that gloss will also handle many other sentences such as "The plate contains punch," etc. Other pertinent sentences for "dish1" are "People serve food on dishes," "Dishes are made of metal, plastic, or ceramic ware," "People put dishes in dishwashers to clean them," etc. All those sentences take about a minute to teach to the system.

4 Algorithm for Determining Noun Senses by Subsumption of Verb Predicates and Semantic Roles

We have designed and implemented several algorithms to resolve noun senses based on the interpreted glosses. Some of the algorithms base their decisions only on the noun senses in the interpreted glosses, while others determine the noun senses of new sentences by using the verb predicate and the semantic roles of the interpreted glosses. For space limitations, we concentrate only on the latter algorithms.

Suppose that we have defined the following gloss for refrigerator1, a monosemous word, "People put food in refrigerators." Then, the system encounters the sentence, "He put the batter in the refrigerator." Let us refer to the sentence being interpreted as IS. As explained, the selectional restrictions for put1 cannot determine the sense of "batter" in the IS. The interpreter's output for the *theme* of the IS (the sentence being interpreted) is : *(obj ((dfart the) (noun batter)) ((ballplayer1 batter1) (concoction1 batter2)) (theme))*. The *theme* of the gloss is: *(obj ((noun food)) ((food food1)) (theme))*. The interpreter cannot decide between the two senses of "batter," and for that reason it prints both senses. The algorithm for noun sense disambiguation obtains all glosses for all noun senses in the IS (the sentence being interpreted). Then, the algorithm (see Figure 3) implemented in Lisp performs the following steps for each GL (gloss). The algorithm will establish the sense of "batter" as batter2, by verifying first that the verb predicate of a GL (gloss) subsumes the verb predicate of the IS (the sentence being interpreted) or that the verb predicate of the IS (the sentence being interpreted) and the verb predicate of the GL (gloss) belong to the same verb predicate hierarchy (a relaxed criterion).

Let LNOUNS-IS be a list containing all noun senses in the IS (the sentence being interpreted). Let L-GL be the list of all interpreted glosses for each of the noun senses in LNOUNS-IS. Let IS-VERB-PRED be the verb predicate of the IS (the sentence being interpreted).
While (L-GL is not empty) do:
Let FIRST-GL be the first gloss in L-GL. If the verb predicate of FIRST-GL subsumes the IS-VERB-PRED or they belong to the same hierarchy, append the output of **Role-Subsumption** (FIRST-GL, IS) to ANSWER. Delete FIRST-GL from L-GL.
End **While**.
Compute FINAL-ANSWER: Obtain the Longest-Sublist in ANSWER (the one with the greatest number of roles). If there is more than one list with the same number of roles in Longest-Sublist, return all noun senses in each role as the answer.
Function Role-Subsumption (FIRST-GL, IS).
Let TEP1, TEP2, TEP3, CORRESP-IS-ROLE, ANSWER, be temporary variables. Let IS-ROLES be a list containing the roles of the IS (the sentence being interpreted). Let ROLES-FIRST-GL be a list containing the roles of the FIRST-GL.
While (ROLES-FIRST-GL is not empty) do:
Let ROLE-GL be the first role in ROLES-FIRST-GL. Search for GL-ROLE in IS-ROLES and if it is found assign it to the variable CORRESP-IS-ROLE. If it is not found, assign NIL to CORRESP-IS-ROLE.
Assign to TEP1 the noun sense of ROLE-GL. Assign to TEP2 the noun senses of CORRESP-IS-ROLE.
Assign to TEP3 all noun senses in TEP2 subsumed by the noun sense in TEP1. Append TEP3 to ANSWER. Delete the first role from ROLES-FIRST-GL. Set TEP3 to NIL.
End **While**
Return ANSWER End **Function Role-Subsumption**

Fig. 3. Algorithm1 - Determining Noun Senses by Subsumption of Verb Predicates and Semantic Roles

If that is not the case, the algorithm discards that GL (gloss) and gets the next GL (gloss) if any. In this example, the predicate of the GL (gloss) subsumes the predicate of the IS (the sentence being interpreted).

Then, the algorithm checks if the noun sense in the semantic roles of the GL (gloss) subsumes the noun sense, or senses, in the semantic roles of the IS (the sentence being interpreted). The comparison is done *agent* of the GL (gloss) against *agent* of the IS (the sentence being interpreted), *theme* of GL (gloss) against *theme* of the IS (the sentence being interpreted), and so on and so forth. There is only a noun sense in the roles of the GL (gloss), because the noun senses of the glosses are resolved, while the roles of the IS (the sentence being interpreted) may have one or more noun senses.

Thus, for the aforementioned example we can verify that: *1)* the noun sense of the *agent* of the GL (gloss) subsumes the noun sense of the *agent* of the IS (the sentence being interpreted), *2)* the noun sense of the *theme* of the GL (gloss), namely *food1*, subsumes one of the noun senses of "batter," namely batter2, in the *theme* of the IS (the sentence being interpreted), and *3)* the noun sense of the *goal* of the GL (gloss) subsumes the noun sense of the *goal* in the IS (the sentence being interpreted).

As a result, the algorithm returns the following: *(refrigerator1 put (agent she) (theme batter2) (goal refrigerator1))*. Had the sentence been "He put the bass/shrimp in the refrigerator," the algorithm would perform likewise for "bass," and "shrimp" (the first sense of "shrimp" in WordNet 1.6 is small person). However, in the case of "bass" the algorithm will output two senses for "bass," bass4 (sea bass) and bass5 (freshwater bass) because both senses have *food* as their hypernym in WN. The algorithm outputs all the roles in the IS (the sentence being interpreted) subsumed by the roles in each GL (gloss).

Simple glosses such as "People put food in a refrigerator1/oven1/pot1/dish1 ..." go a long way in covering many noun senses for which no gloss has been defined. Consider the sentence "In the operation, the appendix was removed." Even if there is no gloss for appendix2 (a body part), the following gloss for operation7 "In an operation, some body parts are removed" will resolve not only the sense of "operation" as operation7 (surgical operation) but also the sense of "appendix" as appendix2. Likewise, the gloss "Sometimes in a demonstration, demonstrators may be hurt," will solve the sense of "demonstrator" as demonstrator3 and "demonstration" as demonstration3. The gloss for cell2, "All living things are made of cells," will resolve not only the sense of "cell" as cell2 in "Plants are made of cells," but also that of "plant" as plant2. The gloss under plant2, "Plant parts become plants," will solve the sense of "plant" as plant2 and "runner" as runner5 in "The runner became a new plant." The examples can be easily multiplied.

Algorithm1-B

This algorithm is a relaxation of algorithm1. In the cases in which the verb predicate of the GL (gloss) and the IS (the sentence being interpreted) do not belong to the same predicate hierarchy, algorithm1-B finds subsumptions only between the semantic roles of the GL (gloss) and the IS (the sentence being

Let Head-Noun-NP be the head noun of the NP and Head-Noun-PP be the head noun of the object of the PP in [NP PP].
Collect all glosses for the noun senses of Head-Noun-PP having a *goal* role, and put them in the variable GLS-GOAL.
While GLS-GOAL do
Let FIRST-GL be the first gloss in GLS-GOAL
If the head noun of the *theme* in FIRST-GL subsumes any of the senses of Head-Noun-NP and the head noun of the *goal* in FIRST-GL subsumes any of the senses of Head-Noun-PP, put all senses subsumed into the variable TEP1. Insert TEP1 into ANSWER. Set TEP1 to NIL.
Delete the first GL (gloss) from GLS-GOAL.
end **While** Return ANSWER.

Fig. 4. Algorithm2 - Resolving Noun Senses in [NP PP] segments by Reasoning with *Theme* and *Goal* Roles

interpreted). Its results are taken in account if at least two semantic roles in the IS (the sentence being interpreted) are subsumed.

5 Algorithm for Resolving Noun Senses in [NP PP] Segments by Reasoning with Theme, Goal and At-Loc Roles

This section explains how to use the interpreted glosses for resolving the senses of the head nouns of prepositional phrases modifying noun phrases. Consider the sentence, "The plants in the pot vanished," or "The batter in the refrigerator vanished." The algorithm explained in the previous section cannot help to determine the sense of "plant," or "batter" in these examples. However, the semantic roles in some of the glosses for pot4 and refrigerator1 can be used to solve the senses of "plant," "pot" and "batter" in a very decisive manner. The algorithm (see Figure 4) works as follows. Suppose that the [NP PP] to be interpreted is "The batter in the refrigerator," and that we have a gloss stored under refrigerator1 that says, "People put food in refrigerators." When a preposition that may stand for the semantic role *goal* (e.g., "into," "onto," "in," "on") follows a NP, the glosses for the noun senses of the head noun of the PP are searched and all those that have a *goal* role are collected. In our example, all glosses for "refrigerator" are searched, and those with a *goal* role are collected. Then for each GL (gloss) with a *goal* role, the algorithm verifies: *a)* if the head noun of the *theme* of the GL (gloss) subsumes any of the senses of the head noun of the NP in the [NP PP] segment ("batter" in our example) and *b)* if the *goal* of the GL (gloss) subsumes any of the senses of the head noun of the object of the PP ("refrigerator" in the example). If both *a)* and *b)* are true, the algorithm returns all noun senses subsumed by the *theme* and *goal* of the gloss. For this example, the output is (ALGORITHM2= (BATTER2 REFRIGERATOR1)).

Another version of this algorithm reasons with *theme* and *at-loc* roles. This algorithm is activated when the preposition in the [NP PP] segment may stand for an *at-loc* role, namely the prepositions “in” “on” and “at,” “near,” “along,” “outside,” etc. Suppose that we want to determine the sense of “bay” in “The ship in the bay.” Let us further assume that we have stored the gloss “Vessels are found in bays” under bay1 (a body of water). This algorithm collects all glosses for the senses of “bay” that have an *at-loc* role, and is identical to the one in Figure 4 if one replaces *goal* with *at-loc*. This algorithm will be able to clearly determine the different senses of “bay” in “the ship in the bay” and “the bay (compartment) in the ship,” given a gloss for bay4 (a ship compartment).

6 Testing

We have tested the algorithms (algorithm1, algorithm1-b, algorithm2) in sentences most of them taken from the (*The World Book Encyclopedia*, World Book, Inc. Chicago) and some from the BNC corpus. First, we looked into the Senseval-2 dataset which contains 29 words. Unfortunately, this data set contains few examples of cases to which the algorithms explained in this paper apply. However, many of the examples in the Senseval-2 dataset can be solved by other algorithms based on our methods, and explained in a forthcoming paper. In that paper, we show how the sense of “conductor” in the sentence “But in the twentieth century, conductors have replaced composers as the most influential people in musical life” (BNC corpus) can be solved by using only the noun senses in the interpreted glosses. The reason why the sense of “conductor” in that sentence cannot be solved by these algorithms is because the verb predicate for “replace” does not determine the sense of “conductor.” The same applies to the sentence “The conductor fired the violinist.” However, the situation is very different if the verb is “direct.”

The purpose of this test has been to show that the algorithms explained in this paper solve noun senses which will be very hard to solve by other algorithms. We have defined and semantically interpreted 1235 glosses for 706 noun senses, or concepts. We have chosen 34 words (*conductor, striker, operation, arm, nail, table, rally, plant, chair, cell, beam, blow, dish, spring, pot, bed, ball, colony, sign, bat demonstrator, plot, crane, port, pen, star, paddle, mast, article, dam pocket, coat, bay, cabin*), and searched for sentences with those words. Some of the sentences were formed by students while we were testing the system. The sentences were parsed and semantically interpreted, and, then, the algorithms were applied. The algorithms were tested in 120 relevant sentences, identifying correctly the target noun sense in 87%. The algorithms were unable to identify the noun sense correctly, or they came up with more than one sense for the target noun in 13% of the 120 sentences. Most of the failures can be corrected by adding some basic glosses to some noun senses. Other problems are due to the semantic interpreter, or implementation. We have built a small file[1] containing 100 sentences for which the algorithms found the correct sense for the target noun. All 34 words are represented in the file. First we list the target word followed by colon, and, then, the test sentence for that word, followed by the

[1] www.cs.ucf.edu/~gomez/nounsenses-disambiguation

output of the algorithms. The file also contains some comments explaining the output. These are two sample sentences for the word “port”:

```
(p "they make port in Portugal from grapes")
(ALGORITHM1=
 ((PORT2 MAKE-OR-CREATE-SOMETHING (THEME PORT2) (OF-STUFF GRAPE1) )))
Comment: port2 is inheriting glosses from wine1. No glosses have been
provided for port2.

(p "they brought the grapes to the port in a boat")
(ALGORITHM1= ((PORT1 BRING-THINGS (THEME GRAPE2 GRAPE1) (GOAL PORT1)
  (INSTRUMENT BOAT1) )))
```

7 Related Work and Discussion

This work falls within the knowledge-based approaches to NSD [8]. Most work on NDS has been based on semantic similarity in the WordNet ontological taxonomy [14,9]. A major distinctive aspect of our work is that the algorithms use interpreted sentences in which verb senses and noun senses are solved. Another differential aspect is that the task of NSD is linked to the overall task of semantic interpretation. The circularity between NSD and selectional restrictions is overcome by giving priority to the verb selectional restrictions, which narrow the noun senses. Then, the NSD algorithms are used to resolve the final senses. Thus, it becomes critical that the verb selectional restrictions do not select incorrect senses, or over-narrow them. As a result, the NSD algorithms are casting light on the task of defining the selectional restrictions for the verb predicates. A work that also uses semantic interpretation for acquisition is [5]. The authors present algorithms for the acquisition of linguistic knowledge and domain knowledge from texts. The learning is realized by a classifier in a terminological representation system. In contrast, in our research the acquisition occurs by being told and aims at acquiring commonsense knowledge that will permit further understanding.

The other aspect of this work is the acquisition of common sense knowledge by being told. Wordnet is a partial realization of Quillian’s dream [12] of building a general semantic network of common sense concepts, because it relates concepts only by *is-a* and *part-of* links. The acquisition methods described in this paper link concepts by semantic relations expressing events, actions. etc. A good interface providing some validation of the definitions will allow students or volunteers to populate WordNet noun senses with the interpreted glosses creating a general common sense knowledge network, that could be used for all kinds of natural language understanding tasks. A major difference between this work and ConceptNet [7] is that the sentences entered in our system are parsed and semantically interpreted by determining their semantic roles and verb predicate. The semantic structures built from the user’s sentences are fully disambiguated, with the main concepts linked to WordNet noun ontology, and our ontology of verb predicates. This rich knowledge structure is what permits the semantic interpreter to bootstrap its knowledge in order to deal with other semantic tasks.

8 Conclusions

We have explained methods to bootstrap the knowledge of a semantic interpreter. The semantic interpreter acquires knowledge about ordinary concepts by being told, and applies this knowledge to noun sense disambiguation. In order to show the relevance of the knowledge acquired, we have designed and tested three algorithms for noun sense disambiguation.

References

1. Fellbaum, C.: English Verbs as a Semantic Net. In: Fellbaum, C. (ed.) WordNet: An electronic Lexical Database and some of its applications, pp. 69–104. MIT Press, Cambridge (1998)
2. Gomez, F.: Building verb predicates: A computational view. In: Proceedings of the 42nd Meeting of the Association for Computational Linguistics, ACL 2004, Barcelona, Spain, pp. 351–358 (2004)
3. Gomez, F.: Semantic Interpretation and the Upper-Level Ontology of WordNet. Journal of Intelligent Systems 16(2), 93–116 (2007)
4. Gomez, F.: An algorithm for aspects of semantic interpretation using an enhanced wordnet. In: Proceedings of the 2nd North American Meeting of the North American Association for Computational Linguistics, pp. 87–94 (2001)
5. Hahn, U., Marko, K.G.: An integrated dual learner for grammars and ontologies. Data & Knowledge Engineering 42, 273–291 (2003)
6. Lesk, M.: Automatic sense disambiguation using machine readable dictionaries: How to tell a pine cone from an ice cream cone. In: Proceedings of the 1986 ACM SIGDOC Conference, Toronto, pp. 24–26 (1986)
7. Liu, H., Singh, P.: ConceptNet - a practical commonsense reasoning tool-kit. BT Technology Journal 22, 211–226 (2004)
8. Mihalcea, R.: Knowledge-based methods. In: Agirre, E., Edmonds., P. (eds.) Word Sense Disambiguation, pp. 107–127. Springer, Heidelberg (2006)
9. Mihalcea, R., Moldovan, D.: A method for word sense disambiguation of unrestricted texts. In: Proceedings of the Annual Meeting of the Association for Computational Linguistics, College Park, Maryland, pp. 152–158 (1999)
10. Miller, G.: Nouns in wordnet. In: Fellbaum, C. (ed.) WordNet: An electronic Lexical Database and some of its applications, pp. 23–46. MIT Press, Cambridge (1998)
11. Moldovan, D., Russ, V.: Logic form transformation of wordnet and its applicability to question answering. In: Proceedings of the 39th meeting of the ACL, Toulouse, France, pp. 402–409 (2001)
12. Quillian, M.: Semantic memory. In: Minsky, M. (ed.) Semantic Information Processing, pp. 216–270. MIT Press, Cambridge, Mass (1968)
13. Quine, V.: Word and Object. MIT Press, Cambridge (1960)
14. Resnik, P.: Semantic similarity in a taxonomy: An information-based measure and its application to problems of ambiguity in natural language. Journal of Artificial Intelligence Research 11, 95–130 (1999)
15. Wittgenstein, L.: Philosophical Investigations. Blackwell, Oxford (1958)