

Semantic Selectional Restrictions for Disambiguating Meronymy Relations

A non-representational DRT-like discourse semantics

Ashwin Ittoo¹, Gosse Bouma²

¹Faculty of Economics and Business,

²Faculty of Arts,

University of Groningen,

Groningen, The Netherlands

Abstract

We present an unsupervised approach to automatically learn lexico-syntactic patterns encoding meronymy relations from texts. Our major contribution lies in alleviating the challenge of disambiguating polysemous patterns that encode meronymy only in some contexts. We rely on the linking theory to posit that semantic features of the Part and Whole instances participating in a meronymy relation facilitate the identification of meronymy-encoding patterns. We abstract the instances to their hypernyms, enforcing semantic selectional restrictions to constrain the contexts within which patterns participate in meronymy. We disambiguate polysemous patterns using their contexts based on a modified version of Harris' distributional hypothesis by postulating that similar patterns share similar contexts. Our experiments revealed that enforcing selectional restrictions enables detecting high-quality patterns. Furthermore, our method does not require annotated data, and has a broader coverage compared to previous studies.

1 Introduction

Meronymy is a semantic relation between an object corresponding to a “part” and to its corresponding “whole” (Girju et al. 2006). If an entity X is the meronym of another entity Y, then sentences of the form “*Xs are parts of Y*” or “*Y has Xs*” are valid when noun phrases X and Y are interpreted generically; for example, “*an engine is a part of a car*”. The inverse of meronymy is holonymy. Meronymy relations between part and whole objects are crucial for various Natural Language Processing tasks, such as question-answering and information extraction (Girju et al. 2003).

Several research efforts, geared towards automatically identifying meronymy patterns from texts, have been proposed (Girju et al. 2006, Girju et al. 2003, Berland and Charniak 1999). However, the long-standing challenge of resolving pattern ambiguity has not yet been adequately addressed. Pattern ambiguity arises when an expression encodes meronymy only within specific contexts. For example, the genitive pattern “*of*” is polysemous since it conveys meronymy in

“*engine of the car*”, but not in “*book of the student*”, in which it indicates a *possession* relation.

In this paper, we present a novel, unsupervised approach to automatically learn high-quality lexico-syntactic patterns encoding meronymy relations from unstructured, natural language texts. We address the pattern ambiguity issue by considering part-whole relations as analogous to semantic frames (Fillmore 1976, Gildea et al. 2002). Semantic frames schematically depict actions or relations, together with the participating concepts that are labeled depending on their roles (Gildea et al. 2002). Based on this analogy, meronymy is a specialized frame, illustrating a part-whole relation between an instance playing the role of a PART, and another one with the corresponding WHOLE role. The instances occurring in the PART and WHOLE roles of a pattern constitute its contexts, and determine, for ambiguous patterns, whether they encode meronymy. The novelty in our approach to disambiguate polysemous meronymy expressions lies in enforcing semantic selectional restrictions on the instances in the PART and WHOLE roles of patterns to constrain the possible contexts within which these patterns encode part-whole relations. Selectional restrictions are obtained by abstracting the PART and WHOLE instances to their hypernyms. We disambiguate polysemous patterns based on a modified version of Harris’ distributional hypothesis (Harris 1968), in which we claim that “patterns are similar to the extent to which they share similar contexts”. Hence, to disambiguate a pattern (i.e. to determine whether it encodes meronymy given its contexts), we find out if its contexts are similar to any of the meronymy-indicative contexts within which the pattern can participate in a part-whole relation. Besides acting as selectional restrictions to facilitate the disambiguation of polysemous expressions, converting PART and WHOLE instances to their conceptual classes also helps in broadening the scope of our approach, and in alleviating issues related to data sparsity. Furthermore, our strategy detects meronym noun pairs that are not documented in the WordNet semantic database (Fellbaum 1998). To enhance the quality of the patterns acquired by our unsupervised procedure, we rely on certain mathematical and logical properties of meronymy relations.

Compared to previous related studies, our methodology has broader coverage, and its unsupervised learning procedure reduces the reliance on large amount of annotated training data. We also innovate in our choice of the English Wikipedia corpus for mining and evaluating meronymy relations. Wikipedia has till now been exploited to extract very general relations (Nguyen et al. 2007, Suchanek et al. 2007).

The remainder of this paper is structured as follows. We present an overview of the theories of parthood in Section 2. Section 3 discusses some recent studies related to the automatic discovery of part-whole relations. We present our unsupervised approach for automatically identifying meronymy relations in Section 4. Our experiment results follow in Section 5. We conclude and highlight areas for future investigation in Section 6.

2 Meronymy Relations

In linguistics, the semantic relation between part entities and their corresponding wholes is known as meronymy (Winston et al. 1987). An entity X is the meronym of another entity Y if “*Xs are parts of Y*”, or “*Y has Xs*”. Conversely, Y is the holonym of X. Meronymy relations can be established both at the conceptual (class) level and at the individual instances (objects) level (Cruse 1986). Class-level part-whole relations indicate that every instance of the WHOLE concept includes one or more instances of the PART concepts. Instance-level meronymy relations imply that the (specific) WHOLE object includes one or more of the PART objects.

2.1 Formal and Mathematical Properties of Parthood

Simons (Simons 1987, Simons 1991) and Varzi (Varzi 2004) defined meronymy as a strict partial ordering, with the following axioms: existence, asymmetry, supplementarity, transitivity, extensionality, irreflexivity, existence of mereological sum, existence of mereological atoms, and decomposition into atoms. Of particular interest to us are asymmetry, transitivity and irreflexivity, which we rely upon to validate the part-whole relations acquired by our approach. They are illustrated in Equations (1) - (3). $PWR(X, Y)$ indicates a part-whole relation between X and Y.

The irreflexivity property states that a concept X cannot be a part or a whole of itself (Keet 2006), and is formalized in equation (1). The irreflexivity property states that a concept X cannot be a part or a whole of itself (Keet 2006), and is formalized in equation (1).

$$\forall X; \neg PWR(X, X) \quad (1)$$

Equation (2) illustrates the asymmetry property, according to which, if X is part of Y, then Y is not part of X (Keet 2006).

$$\forall X, Y; PWR(X, Y) \rightarrow \neg PWR(Y, X) \quad (2)$$

The transitivity property, in equation (3), states that if X is part of Y, and Y is part of Z, then X is also part of Z (Keet 2006).

$$\forall (X, Y, Z); PWR(X, Y) \wedge PWR(Y, Z) \rightarrow PWR(X, Z) \quad (3)$$

Although transitivity is a minimal requirement for part-whole relations (Artale et al. 1996), it applies only across relations of the same types according to Winston’s taxonomy (Winston et al. 1987).

2.2 Taxonomy of Meronymy Relations

Winston et al. (Winston et al. 1987) developed a taxonomy of part-whole relations based on psycho-linguistic experiments, and classified six major types of

Ambiguous Genitives	ENGINE “of the” CAR BOOK “of the” STUDENT	True False
Ambiguous Noun Compounds	CAR ENGINE STUDENT BOOK	True False
Ambiguous Prepositional Constructs	CAR “with” ENGINE STUDENT “with” BOOK	True False

Table 1: Ambiguous Meronymy Patterns.

meronymy relations. They are: 1) Component-Integral, the relation between a component (part) and its integral (whole), as in ENGINE-CAR, 2) Member-Collection, membership in a collection, such as SHIP-FLEET, 3) Portion-Mass, the relations between portions and masses, and physical dimensions, for example, METRE-KILOMETRE or SLICE-PIE, 4) Material-Object, the relation between an object and its constituent material, such as GRAPE-WINE, 5) Place-Area, the relation between areas and locations within them, as in NETHERLANDS-EUROPE, and 6) Feature-Activity, the relation between phases of an activity, for example, SCORING-PLAYING.

2.3 Ambiguity of Meronymy Encoding Patterns

The many different ways in which an entity can be expressed as a part of another entity give rise to a wide variety of lexico-syntactic structures that can encode meronymy (Iris et al. 1988). Previous studies (Girju et al. 2006) recognized four types of meronymy-indicating lexico-syntactic patterns. They are: 1) genitives and the verb “to have”, 2) noun compounds, 3) prepositional phrases, 4) other rare patterns, (e.g.: “*is a branch of*”).

Some lexico-syntactic patterns are unambiguous and always depict meronymy. Other polysemous patterns are ambiguous, and convey meronymy depending on the contexts in which they occur. Table 1 lists examples of ambiguous meronymy patterns for genitives, noun compounds and prepositional phrases. “True” indicates that the ambiguous pattern encodes meronymy, while “False” indicates the opposite.

Disambiguating polysemous lexico-syntactic structures, to determine whether their occurrences are indicative of meronymy, is a challenge that has not been adequately addressed in past research.

3 Related Work

Hearst (Hearst 1998) proposed an approach for identifying meronymy relations from text based on the occurrence of a restricted set of lexico-syntactic patterns, following a similar successful technique for detecting hypernymy (Hearst 1992). However, the performance in meronymy identification was low, possibly due to the wider variety of meronymy-invoking constructs. Berland and Charniak (Berland

and Charniak 1999) also relied on specific lexico-syntactic patterns to discover part-whole relations. They focused on a small set of genitive patterns and on six seed nouns representing WHOLE concepts. Statistical measures over a large corpus were then employed to locate meronymy relations.

Neither Hearst nor Berland and Charniak discriminated between ambiguous and unambiguous meronymy constructs, which resulted in the low performance of their approaches. They also suffered from low coverage, compounded by their small seed sets, and their discarding of terms with suffixes “ing”, “ness”, or “ity”. Girju et al. (Girju et al. 2006, Girju et al. 2003) attempted to alleviate the pattern ambiguity challenge by machine-learning techniques based on decision tree classifiers. Despite being more accurate than (Berland and Charniak 1999) and (Hearst 1998), they required substantial manual intervention, for example, to identify an initial set of meronymy-encoding patterns, to annotate positive and negative examples, and to disambiguate word senses. Our methodology to automatically acquire high-quality meronymy lexico-syntactic patterns addresses the above-mentioned issues, as well as the crucial challenge of resolving meronymy pattern ambiguity.

4 Methodology

The overall framework in which we embed our unsupervised algorithms to automatically discover meronymy relations consists of a knowledge acquisition and an evaluation phase. Similar to other data-driven methodologies, we start by acquiring seed PART-WHOLE noun pairs from a semantic database. These *noun* pairs are abstracted to their corresponding PART-WHOLE *concept* pairs. Lexico-syntactic patterns relating (noun) instances of the PART-WHOLE concept pairs are then searched for, validated and extracted from a corpus. Evaluation involves measuring the precision with which the acquired patterns encode meronymy over a completely different test corpus.

4.1 Acquiring Seeds

The aim of this phase is to automatically acquire meronym or holonym noun pairs, $\langle n_1, n_2 \rangle$, participating in part-whole relations. This means either n_1 is a meronym of n_2 (n_2 is the holonym of n_1) or vice-versa. The pair n_1 and n_2 constitute a PART-WHOLE noun pair. We implemented a meronym/holonym noun pair harvesting procedure that takes as input a seed noun, s_i , and traverses WordNet to collect its meronyms and holonyms for all its different senses. The collected meronyms and holonyms are then treated as seeds, and fed to our harvesting algorithm so that their meronyms and holonyms are acquired recursively. Unlike past studies (Berland and Charniak 1999, Girju et al. 2003, Girju et al. 2006, Hearst 1998), our seed harvesting and expansion strategy enables us to broaden our coverage since the holonym of a seed s_i could have other meronyms s_j (besides s_i). For example, $s_i = \text{“engine\#3”}$, has as holonym “train\#1” , which in turn has another meronym synset $s_j = \langle \text{car\#2}, \text{railcar\#1}, \text{railway_car\#1},$

railroad_car#1 >. Our approach also efficiently handles seeds that are both meronyms and holonyms, such as $s_i = \text{“engine\#3”}$ which is a holonym of *“footplate\#1”*, but a meronym of *“train\#1”*. Furthermore, our seed collection procedure requires the specification of either a part or a whole noun while previous studies needed both the part and whole seeds.

Thus, from a minimal list of input nouns, this phase outputs an augmented seed set of PART-WHOLE noun pairs. Since they were harvested by traversing the meronymy and holonymy semantic links of WordNet, which we consider a reliable reference, the co-occurrences of these noun pairs in any lexico-syntactic patterns always indicate meronymy.

4.2 Conceptual Semantic Abstraction

This stage takes as input the previously harvested PART-WHOLE *noun* pairs, and generalizes them to their super-ordinate concepts (classes), resulting into PART-WHOLE *concept* pairs. We achieve this abstraction by replacing the PART and WHOLE noun instances with their respective direct ancestors, one level up the WordNet hypernymy chain. For example, the PART-WHOLE *noun* pair $\langle \textit{grape}, \textit{wine} \rangle$ will be abstracted to the PART-WHOLE *concept* pair $\langle \textit{edible_fruit\#1}, \textit{alcohol\#1} \rangle$, where $\#x$ indicates WordNet sense x .

Our rationale for incorporating semantic features such as hypernyms stems from the linking theory (Fillmore 1976, Gildea et al. 2002), which is at the core of frame-based semantic role labeling (Gildea et al. 2002). According to this theory, lexico-syntactic realizations of patterns that relate arguments to their predicates can be accurately predicted from the arguments’ semantics. Our reformulation of meronymy relation discovery as a special case of semantic role labeling, with meronymy relations corresponding to semantic frames, enables us to postulate that semantic information about PART and WHOLE arguments could facilitate identifying lexico-syntactic realizations of patterns that encode meronymy (Section 4.3). Abstracting PART and WHOLE noun pairs to their hypernyms also enables the implicit capture of certain mathematical properties (irreflexivity, asymmetry, transitivity) as well as other characteristics (homogeneous and homeomeric) of the relations in which they participate. Furthermore, replacing individual noun instances with their conceptual classes broadens the coverage of lexical items that are characterized by large vocabularies, and hence, reduces problems of data sparsity. The output of this phase is a set of PART-WHOLE concept pairs of the form $\langle \textit{Hypernym}(n_1\#x), \textit{Hypernym}(n_2\#y) \rangle$, where n_1 and n_2 are PART-WHOLE noun pairs with WordNet sense numbers x and y respectively. The function $\textit{Hypernym}(w\#s)$ returns the immediate hypernym of noun w with sense s based on WordNet’s hierarchy. When they appear as arguments of lexico-syntactic patterns, these concept pairs impose semantic selectional restrictions that constrain the possible contexts within which these pattern encode meronymy, facilitating the disambiguation of polysemous patterns (Section 4.3). We index such PART-WHOLE concept pair in a PART-WHOLE concept lexicon.

One issue with generalizing noun instances to their immediate WordNet hy-

pernyms is that the most informative hypernym could reside a few level up the WordNet hierarchy. For example, abstracting the noun pair $\langle robin, wing \rangle$ to its corresponding concept pair yields $\langle thrush, organ \rangle$, while the pair $\langle bird, organ \rangle$ would have been more desirable. We discuss how to handle this problem in our future work (Section 6).

4.3 Acquiring and Disambiguating Meronymy-Indicating Patterns

The objective in this step is to acquire lexico-syntactic expressions that convey meronymy from a corpus. This process entails the proper disambiguation of polysemous patterns to determine whether their occurrences indicate meronymy. For each identified PART-WHOLE concept pairs $\langle Part_Concept, Whole_Concept \rangle$ from the previous phase, we search a portion of the English Wikipedia corpus for sentences containing noun pairs $\langle n_1, n_2 \rangle$ such that n_1 is an instance of the PART concept *Part_Concept*, and n_2 is an instance of the corresponding WHOLE concept *Whole_Concept*.

Since *Part_Concept* and *Whole_Concept* participate in meronymy at the class-level, according to Cruse (Cruse 1986), every instance of type *Whole_Concept* (e.g.: n_2) should contain one or more instances of type *Part_Concept* (e.g.: n_1). Based on Cruse's idea and on the extensionality property of part-whole relations, which states that objects with the same parts are identical, we infer that n_1 and n_2 are PART-WHOLE noun pairs that also participate in meronymy relations. As example, consider the PART-WHOLE concept pair $\langle edible_fruit\#1, alcohol\#1 \rangle$ from the PART-WHOLE concept lexicon (Section 4.2). Searching for instances of this meronymy concept pair in the Wikipedia corpus could lead to the identification of the PART-WHOLE noun pair $\langle grape\#1, wine\#1 \rangle$ from Wikipedia sentences. (We do not discuss determining the correct senses of nouns in this paper). Since these identified noun pairs are instances of known PART-WHOLE concept pairs, and hence also participate in part-whole relations, the lexico-patterns that relate them in their sentences encode meronymy. Similar to (Snow et al. 2005), we formalize the representational space of lexico-syntactic patterns by dependency paths. We syntactically parse sentences containing occurrences of the PART and WHOLE noun pairs, and define a meronymy lexico-syntactic construct as the shortest path in the parsed dependency structure that relates the PART and the WHOLE nouns. For example, given sentence

S1= “Plan Bordeaux calls for a simplification of French wine labels , including the name of the grapes ...”.

S1 was identified from our corpus since it contains instances of the PART-WHOLE concept pair $\langle edible_fruit\#1, alcohol\#1 \rangle$, namely the PART and the WHOLE noun pairs “grape” and “wine”. The shortest dependency path between the PART and the WHOLE noun pairs extracted from it is

P= “+nn+labels+pobj+of<- prep ->including+pobj+name+prep+of+pobj+”.

In the above example, the lexico-syntactic pattern *P* encodes meronymy since its arguments are instances of the known PART-WHOLE concept pair “*edible_fruit*” and “*alcohol*”. In this way, we acquire a set of promising meronymy lexico-

syntactic patterns by searching our corpus for sentences containing instances of previously identified PART-WHOLE concept pairs, parsing these sentences and extracting the shortest dependency paths between the noun instances. We index the learnt patterns in a meronymy-pattern lexicon.

However, our approach could prove to be too general, and identify instances that do not participate in meronymy relations, such as $\langle kiwi\#1, wine\#1 \rangle$ which are valid instances of $\langle edible_fruit\#1, alcohol\#1 \rangle$. Subsequently extracted patterns relating such false meronymy noun pairs do not encode part-whole relations. To keep only those constructs that are likelier to express meronymy, we maintain a frequency count, and discard patterns with frequencies below an experimentally set threshold. We assume that true meronymy patterns such as “*made of*” will not relate instances that are not parts and wholes, as in “*wine is made of kiwi*”. Furthermore, invalid meronym noun pairs are in most cases related by conjunctions or negations, as in “*he likes wine and kiwi*” or “*wine is not made of kiwi*”, which we filter in our approach to improve the quality of the acquired part-whole patterns.

To disambiguate polysemous patterns, such as genitives (“*of*”) that encode meronymy only in certain contexts, we rely on the previously acquired PART-WHOLE concept pairs indexed in the PART-WHOLE concept lexicon, and on the unambiguous patterns recorded in the meronymy-pattern lexicon. The key idea underlying our disambiguation procedure is that PART-WHOLE concept pairs enforce semantic selectional restrictions on the lexico-syntactic patterns they co-occur with, constraining the possible contexts in which these patterns express meronymy, thereby facilitating their disambiguation. Contextual information has been traditionally employed in word sense disambiguation, following Harris’ distributional hypothesis, which states that words are similar to the extent to which they share similar contexts. We reformulate this hypothesis as “lexico-syntactic patterns are similar to the extent to which they share similar contexts”, and use this modified hypothesis in disambiguating polysemous patterns. In our case, the contexts of patterns are defined by the arguments that they sub-categorize. For example, given an unambiguous meronymy lexico-syntactic pattern $P = \text{“contains”}$, acquired from our corpus and indexed in the meronymy-pattern lexicon, and its meronymy-indicating contexts consisting of the PART-WHOLE concept pair $\langle vehicle\#1, motor\#1 \rangle$. P expresses a conceptual level part-whole relation “*vehicle contains motor*”. To determine whether the occurrence of an ambiguous pattern P' (e.g.: “*of*”) with arguments “*args1*” (e.g.: “*engine*”) and “*args2*” (e.g.: “*car*”) encodes meronymy, we compute its contextual similarity with the meronymy-pattern P . We consider P' to be similar to P , and hence, to also encode a part-whole relation, if the arguments defining its contexts are instances of the known PART-WHOLE concept pairs that are sub-categorized by the meronymy pattern P . Thus, in this example, since “*engine*” and “*car*” are instances of “*motor\#1*” and “*vehicle\#1*” respectively, we infer that the ambiguous meronymy pattern $P' = \text{“of”}$ indicates a part-whole relation. The semantic classes (e.g.: concepts “*motor\#1*” and “*vehicle\#1*”) act as selectional modifiers on ambiguous patterns, and restrict the allowable contexts in which these patterns participate in meronymy.

Had the contexts of P' been characterized by the arguments “*book*” and “*student*”, as in “*book of student*”, our technique would infer that P' does not encode meronymy in these contexts. It reaches this conclusion on two bases. First, the corpus from which the patterns are extracted does not contain invalid English constructs involving meronymy patterns as in “*student is made of book*”. Such invalid expressions will cause our algorithm to identify “*student*” and “*book*” as PART-WHOLE noun pairs, and to infer that the pattern “*made of*” encodes meronymy when it sub-categorizes instances of type “*enrollee#1*” and “*publication#1*” (the conceptual classes of “*student*” and “*book*” respectively). The second basis is that meronymy and holonymy relations documented in WordNet are correct. Otherwise, we will count pairs such as $\langle \text{“enrollee\#1”}, \text{“publication\#1”} \rangle$ as valid PART-WHOLE concepts, and will wrongly consider the patterns in which their instances occur, such as “*student buys book*” or “*student has book*”, as encoding meronymy.

Our approach, although simple and intuitive, enables the discovery of meronymy (and holonymy) relations between pairs that are not mentioned in WordNet. For example, consider the PART-WHOLE noun pair $\langle \text{base\#2}, \text{construction\#4} \rangle$, harvested during our seed acquisition procedure (Section 4.1), and its corresponding PART-WHOLE concept $\langle \text{support\#7}, \text{artifact\#1} \rangle$ (Section 4.2). Our algorithm mines for meronymy lexico-syntactic patterns by searching our corpus for sentences containing co-occurring instances of concepts “*support#7*” and “*artifact#1*”. One such sentence could be “*buttress of excavation*” since the noun pair $\langle \text{buttress}, \text{excavation} \rangle$ is an instance of the PART-WHOLE concept pair $\langle \text{support\#7}, \text{artifact\#1} \rangle$. Besides concluding that the genitive “*of*” expresses meronymy in the context of $\langle \text{support\#7}, \text{artifact\#1} \rangle$, we also deduce that *buttress* is a meronym of *excavation* - a valid fact that is not mentioned in WordNet. Hence, we augment WordNet with new relations to improve its completeness. To further enhance the performance of our unsupervised learning process, we discard relations that do not satisfy the asymmetry and irreflexivity properties of meronymy; for example, “*engine is part of engine*”, which violates irreflexivity or “*car is part of engine*”, which violates asymmetry.

The output of this stage is a set of (both ambiguous and unambiguous) lexico-syntactic patterns, together with the contexts in which they encode meronymy. These meronymy-indicating contexts are PART-WHOLE concept pairs. They enforce semantic selectional constraints on the patterns, and restrict the set of possible contexts in which the patterns participate in meronymy, thereby enabling their disambiguation. This phase also augments the PART-WHOLE concept lexicon (Section 4.2) with new PART-WHOLE concept pairs.

4.4 Evaluating the Acquired Patterns

The objective in this phase is to evaluate the quality of the patterns acquired. We want to determine whether enforcing semantic selectional restrictions on the PART and WHOLE instances of the patterns contributed in disambiguating polysemous

Winston's Relation	Seed	PART-WHOLE noun pairs collected
Component-Integral	engine	< <i>camshaft, engine</i> >, < <i>arrester, attackaircraftcarrier</i> >, ...
Member-Collection	ship	< <i>bay, ship</i> >, < <i>commode, lavatory</i> >, < <i>davit, ship</i> >, ...
Portion-Mass	metre	< <i>metre, decameter</i> >, < <i>angstrom, nanometer</i> >, < <i>adenine, dna</i> >, ...
Material-Object	grape	< <i>grape, grapevine</i> >, < <i>dimocarpus.longan, genus_dimocarpus</i> >, ...
Place-Area	USA	< <i>last_frontier, united_states</i> >, < <i>empire_state, united_states</i> >, ...

Table 2: Some harvested Part-Whole noun pairs (sense numbers omitted).

meronymy patterns. Evaluation involves searching a test corpus for occurrences of the acquired patterns, and measuring the precision with which they are indicative of part-whole relations. As will be shown in the next section, our unsupervised methodology is able to identify highly-precise meronymy lexico-syntactic patterns.

5 Experimental Evaluations

Our meronym noun pair acquisition starts with five seeds that participate in the different types of Winston's part-whole relations (we did not consider Feature-Activity relations, which are realized by verb entailments). These seeds are fed to our harvesting procedure which traverses WordNet, collects their meronyms and holonyms, and recursively expands the seed set. This strategy led to the discovery of some interesting PART-WHOLE noun pairs for a given input seed as shown in Table 2. A total of 12389 PART-WHOLE noun pairs were thus collected, many of them duplicates.

The collected PART-WHOLE noun pairs were abstracted to their corresponding PART-WHOLE concept pairs based on the WordNet hypernymy hierarchy. Duplicate pairs and those violating the irreflexivity and asymmetry properties of meronymy were discarded, resulting in a concise and informative set of 580 PART-WHOLE concept pairs, such as < *military_academy, agency* >, < *letter, bicameral_script* >, < *propeller, internal - combustion_engine* >, < *stroke, table_game* >. These PART-WHOLE concept pairs enforce semantic selectional restrictions on lexico-syntactic patterns by constraining the possible contexts within which these patterns encode meronymy.

Our candidate meronymy pattern acquisition procedure begins by syntactically parsing around 25% of the English Wikipedia corpus (\approx 140M words and 10M sentences), provided by the University of Amsterdam (<http://ilps.science.uva.nl/WikiXML/> n.d.), using the Stanford parser (http://nlp.stanford.edu/software/lex_parser.shtml n.d.). We saved both

the phrase structure tree and the results of the dependency analysis, from which we extracted meronymy-indicating patterns as the shortest dependency paths between instances of our previously acquired PART-WHOLE concepts. This procedure yielded 11351 individual lexico-syntactic expressions together with their co-occurring PART-WHOLE pairs that restrict the contexts in which these expression convey meronymy.

The most frequent meronymy-indicating pattern was:

“+pobj+with+prep+shoulders+nsubj< -led- >prep+in+pobj+parts+prep+of+pobj+” which co-occurred with instances of types “homo#n#2” (such as “world#n#8”) and instances of type “group#n#1” (such as “people#n#1”) 355 times¹. The next most frequent pattern was the ambiguous genitive “of”. It occurred 91 times within different meronymy-indicating concept pairs, such as < feature#n#2, external_body_part#n#1 >, for example in “temple of the head”, which indeed conveys meronymy. This corroborates with the findings of (Girju et al. 2006), who also observed that genitives were the most frequent (and most ambiguous) meronymy patterns. Around 95% of all the acquired patterns had a frequency of one. These typically were domain-specific meronymy patterns.

We evaluated the patterns, P , thus acquired by measuring their precision in expressing meronymy when their contexts are constrained by (i.e.: they sub-categorize) certain PART-WHOLE concept pairs. We parsed an additional 130K sentences from the Wikipedia corpus and extracted 88219 lexico-syntactic patterns together with their co-occurring noun pairs as our test set. For each pattern P and its PART-WHOLE concept pairs, < Part_Concept, Whole_Concept >, we searched for its occurrences, P_{test} , in the test set. Only those P_{test} co-occurring with nouns m_1 and m_2 that are respectively instances of PART_CONCEPT and WHOLE_CONCEPT were considered. If m_1 and m_2 are meronyms or holonyms pairs, we incremented a count $P.true_positive$ to reflect that the lexico-syntactic pattern P indicates meronymy when it occurs within restricted contexts consisting of instances of PART_CONCEPT and WHOLE_CONCEPT. If m_1 and m_2 are not meronyms or holonyms, we incremented a count $P.false_positive$, indicating that P may not always indicate meronymy within the restricted contexts of instances of PART_CONCEPT and WHOLE_CONCEPT. We manually determined whether m_1 and m_2 are meronyms or holonyms instead of relying (automatically) on WordNet since, as shown before, WordNet is sparse and does not document many valid meronymy or holonymy relations that our approach discovers. The precision of a pattern P occurring within the context c (defined by a PART-WHOLE concept pair) is

$$Precision(P_c) = \frac{P.true_positive}{P.true_positive + P.false_positive} \tag{4}$$

We do not compute recall as we are unaware beforehand of the actual num-

¹We later found that the high count for this pattern was due to its repetitive occurrence within a single segment of the corpus, illustrating a case of Wikipedia spamming.

ber of meronymy patterns in the corpus. Table 3 illustrates evaluating the precision with which the genitive pattern “of” encodes meronymy when its contexts are constrained by instances of specific PART-WHOLE conceptual classes. Noun pairs marked with “*” are meronyms that our technique identified but that are not documented in WordNet, while those marked with “**” are false positives, i.e. we wrongly identified them as possible contexts within which “of” conveys meronymy. The precision of all the evaluated patterns is given in Table 4.

PART: <i>Person#n#1</i> WHOLE: <i>People#n#1</i> Example: (<i>warrior#n#1, nation#n#2</i>)*; (<i>inhabitant#n#1, world#n#5</i>)*; (<i>leader#n#1, business#n#8</i>)*; (<i>owner#n#2, land#n#8</i>)*; (<i>creditor#n#1 land#n#8</i>)** Total Count:17, $Precision = \frac{16}{17} = 0.94$
PART: <i>Point#n#6</i> WHOLE: <i>Time Interval#n#1</i> Example: (<i>beginning#n#2, period#n#2</i>)*; (<i>start#n#2, round#n#2</i>)*; (<i>start#n#2, period#n#2</i>)*; (<i>end#n#2, period#n#2</i>)* Total Count:4, $Precision = \frac{4}{4} = 1$

Table 3: Evaluating the precision of the genitive pattern “of” in the test set.

Our results indicate that in 88% of cases, the patterns learnt together with their respective PART-WHOLE concept pairs, which restrict the contexts within which they participate in part-whole relations, expressed meronymy with 100% precision. As example, the pattern “of”, which as indicated in Table4, encodes meronymy when its arguments are instances of the semantic classes “*Point#n#6*” and “*Time Interval#n#1*”, expresses a part-whole relation in “*start of period*”. The high precision obtained validates our underlying hypothesis that semantic features of the PART and WHOLE instances (in the form of their respective PART-WHOLE conceptual classes) participating in meronymy relations facilitate the identification and disambiguation of these relations by constraining the set of allowable contexts within which they encode meronymy. To confirm these findings, we repeated our experiments, but without abstracting the PART and WHOLE noun pairs (Section 4.1) to their semantic classes. We search for the exact occurrences of these noun pairs in the Wikipedia corpus, and extracted potential meronymy-encoding lexico-patterns that relate these noun pairs (Section 4.3). As could be expected, precision was higher, around 90%, but our coverage was much less. The smaller relative gain in precision compared to the much substantial reduction in coverage, did not outweigh the benefits gained by exploiting the semantic features of PART and WHOLE noun pairs. Furthermore, defining meronymy at the class (concept)-level allowed the discovery of meronym noun pairs that are not defined in WordNet, such as (*warrior#n#1, nation#n#2*),(*room#n#1, level#n#8*) . We also observed that semantic selectional restrictions do not suffice to constrain some contexts within which patterns participate in meronymy. One such context is defined by the PART-WHOLE concept pair *Person#1* and *People#1* . Ambiguous patterns occurring with instances of these concepts might not always express meronymy, such as in “*creditor of land*”.

Lexico-Syntactic Pattern	Part Whole Contexts	Precision
"pobj+of+prep" ("of")	Part#n#2 Object#n#1	1
	Room#n#1 Housing#n#1	1
	Administrative_District#n#1 AD#n#1	1
	Object#n#1 Object#n#1	1
	Point#n#6 Time_Interval#n#1	1
	Concept#n#1 Concept#n#1	0.94
	Tract#n#1 Structure#n#1	1
	Room#n#1 Dwelling#n#1	1
	Person#n#1 People#n#1	0.94
	Room#n#1 Structure#n#1	1
	Area#n#6 Document#n#1	1
Facility#n#1 Store#n#2	1	
"pobj+in+prep" ("in")	Structure#n#1 Structure#n#1	1
	Room#n#1 Structure#n#1	0.5
	Administrative_District#n#1 AD#n#1	1
	Time_Unit#n#1 Time_Unit#n#1	0.67
	Region#n#3 Administrative_District#n#1	1
	Area#n#5 Structure#n#1	0.75
	Body#n#1 Organism#n#1	1
	Collection#n#1 Natural_Object#n#1	1
Tract#n#1 Tract#n#1	0.5	
"pobj+in+prep+located+ "partmod" ("located in")	Administrative_District#n#1 AD#n#1	1
	Structure#n#1 Structure#n#1	1
"pobj+as+prep" ("as")	Room#n#1 Structure#n#1	1
	Object#n#1 Object#n#1	0.5
"pobj+at+prep" ("at")	Structure#n#1 Structure#n#1	1
	Area#n#5 Structure#n#1	1
"pobj+by+prep" ("by")	Region#n#3 Administrative_District#n#1	1
"pobj+for+prep" ("for")	Administrative_District#n#1 AD#n#1	1
"pobj+on+prep" ("on")	Structure#n#1 Structure#n#1	1
	Area#n#6 Document#n#1	1
	Geological_Formation#n#1 GF#n#1	1
"pobj+with+prep" ("with")	Body#n#2 Educational_Institution#n#1	1

Table 4: Precision of test patterns in encoding meronymy when constrained by respective contexts.

6 Conclusion and Future Work

We presented an unsupervised approach to automatically acquire high-quality meronymy patterns from texts. Our approach considers meronymy as a specialized semantic frame, depicting a part-whole relation between instances playing a PART and a WHOLE role. To disambiguate polysemous patterns that encode meronymy only within certain contexts, we rely on the linking theory to posit that semantic features of the PART-WHOLE instances could facilitate the identification of lexico-syntactic patterns encoding meronymy. We abstract instances to their hypernyms, enforcing semantic selectional restrictions to constrain the contexts within which the patterns participate in meronymy. These meronymy-encoding contexts are the crux of our disambiguation procedure, in which we extend Harris' hypothesis to postulate that similar patterns share similar contexts. To validate our approach, we measured the precision of the acquired patterns in expressing meronymy over a test corpus. Our evaluation results indicate that 88% of the learnt patterns expressed meronymy with 100% precision when constraints were applied to restrict the contexts within which they participate in part-whole relations. This highlights the major contribution of our methodology in the identification and disambiguation of meronymy-indicating patterns. Furthermore, our unsupervised approach circumvents the need for annotated training data and manual intervention, as opposed to those based on supervised learning. Also, our definition of meronymy at the conceptual level enabled the identification of part-whole noun pairs that are not documented in WordNet as participating in meronymy. By expanding an initial seed set, the technique we present does not suffer from low coverage, as do previous related studies. As future work, we will address the issue of obtaining the *basic level category* (Izquierdo et al. 2007), i.e. most informative hypernym (semantic class) of nouns, instead of abstracting them to their immediate WordNet hypernym. Another improvement could be reducing the reliance on WordNet (or basic level categories) in abstracting nouns to their concepts by using Wikipedia instead. We also intend to validate our findings by calculating the inter-annotator agreement score between judges who have to decide whether a given pattern encodes meronymy when appearing within a particular context.

Acknowledgments.

Ashwin Ittoo is a PhD candidate under the IOP-IPCR program of the Dutch Ministry of Economic Affairs.

References

- Artale, Alessandro, Enrico Franconi, Nicola Guarino, and Luca Pazzi (1996), Part-whole relations in object-centered systems: An overview.
- Berland, Matthew and Eugene Charniak (1999), Finding parts in very large corpora, *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*, Association for Computational Linguistics, Morristown, NJ, USA, pp. 57–64.

- Cruse, D.A. (1986), *Lexical Semantics*, Cambridge University Press, Cambridge, UK.
- Fellbaum, Christiane, editor (1998), *WordNet An Electronic Lexical Database*, The MIT Press, Cambridge, MA ; London. <http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=8106>.
- Fillmore, C.J. (1976), Frame semantics and the nature of language, *Annals of the New York Academy of Sciences* **280**, pp. 20–32.
- Gildea, Daniel, Daniel Jurafsky, Daniel Gildea, and Daniel Jurafsky (2002), Automatic labeling of semantic roles, *Computational Linguistics* **28**, pp. 245–288.
- Girju, Roxana, Adriana Badulescu, and Dan Moldovan (2003), Learning semantic constraints for the automatic discovery of part-whole relations, *NAACL '03: Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology*, Association for Computational Linguistics, Morristown, NJ, USA, pp. 1–8.
- Girju, Roxana, Adriana Badulescu, and Dan Moldovan (2006), Automatic discovery of part-whole relations, *Comput. Linguist.* **32** (1), pp. 83–135, MIT Press, Cambridge, MA, USA.
- Harris, Z. S. (1968), *Mathematical Structures of Language*, Wiley, New York, NY, USA.
- Hearst, M. (1998), Automated discovery of wordnet relations, in Fellbaum, Christiane, editor, *WordNet: An Electronic Lexical Database and Some of its Applications*, MIT Press.
- Hearst, Marti (1992), Automatic acquisition of hyponyms from large text corpora, *Proceedings of the 14th International Conference on Computational Linguistics*, Nantes, France.
- <http://ilps.science.uva.nl/WikiXML/> (n.d.), English wikipedia corpus.
- http://nlp.stanford.edu/software/lex_parser.shtml (n.d.), Stanford statistical parser.
- Iris, M. A., B. E. Litowitz, and M. Evens (1988), Problems of the part-whole relation, in Evens, M. W., editor, *Relational Models of the Lexicon: Representing Knowledge in Semantic Networks*, Cambridge University Press, Cambridge, pp. 261–288.
- Izquierdo, R., A. Surez, and G. Rigau (2007), Exploring the automatic selection of basic level concepts., *Proceedings of the International Conference on Recent Advances on Natural Language Processing (RANLP'07)*.
- Keet, M. (2006), Introduction to part-whole relations: mereology, conceptual modelling and mathematical aspects, *Technical report*, Free University of Bozen-Bolzano, Italy.
- Nguyen, D. P. T., Y. Matsuo, and M. Ishizuka (2007), Exploiting Syntactic and Semantic Information for Relation Extraction from Wikipedia, *IJCAI Workshop on Text-Mining and Link-Analysis*. <http://www.miv.t.u-tokyo.ac.jp/papers/dat-IJCAI07-TextLinkWS.pdf>.
- Simons, P. (1987), *Parts: A Study in Ontology*, Routledge.
- Simons, P. (1991), Part/whole ii: Mereology since 1900, in Burkhardt H., Smith B,

- editor, *Handbook of Metaphysics and Ontology*, pp. 261–288.
- Snow, R., D. Jurafsky, and A. Ng (2005), Learning syntactic patterns for automatic hypernym discovery, *Proceedings of the 17th Conference on Advances in Neural*.
- Suchanek, Fabian, Gjergji Kasneci, and Gerhard Weikum (2007), Yago a core of semantic knowledge, *Proc. Int. Conf. on World Wide Web*, pp. 697–706.
- Varzi, A.C. (2004), Mereology, in Zalta, E.N., editor, *The Stanford Encyclopedia of Philosophy*.
- Winston, Morton E., Roger Chaffin, and Douglas Herrmann (1987), A taxonomy of part-whole relations., *Cognitive Science* **11** (4), pp. 417–444.