A Model for Composing Semantic Relations

Eduardo Blanco and Dan Moldovan
Human Language Technology Research Institute
The University of Texas at Dallas
{eduardo,moldovan}@hlt.utdallas.edu

Abstract

This paper presents a model to compose semantic relations. The model is independent of any particular set of relations and uses an extended definition for semantic relations. This extended definition includes restrictions on the domain and range of relations and utilizes semantic primitives to characterize them. Primitives capture elementary properties between the arguments of a relation. An algebra for composing semantic primitives is used to automatically identify the resulting relation of composing a pair of compatible relations. Inference axioms are obtained. Axioms take as input a pair of semantic relations and output a new, previously ignored relation. The usefulness of this proposed model is shown using PropBank relations. Eight inference axioms are obtained and their accuracy and productivity are evaluated. The model offers an unsupervised way of accurately extracting additional semantics from text.

1 Introduction

Semantic representation of text is an important step toward text understanding, performing inferences and reasoning. Potentially, it could dramatically improve the performance of several Natural Language Processing applications.

Semantic relations have been studied in linguistics for decades. They are unidirectional underlying connections between concepts. For example, the sentence *The construction slowed down the traffic* encodes a CAUSE and detecting it would help answer the question *Why is traffic slower?*

In Computational Linguistics, there have been several proposals to detect semantic relations. Current approaches focus on a particular set of relations and given a text they output relations. There have been competitions aiming at detecting semantic roles (i.e., relations between a verb and its arguments) (Carreras and Màrquez, 2005), and between nominals (Girju et al., 2007; Hendrickx et al., 2009).

In this paper, we propose a model to compose semantic relations to extract previously ignored relations. The model allows us to automatically obtain inference axioms given a set of relations and is not coupled to any particular set. Axioms take as their input semantic relations and yield a new semantic relation as their conclusion.

Consider the sentence *John went to the shop to buy flowers*. Figure 1 shows semantic role annotation with solid arrows. By composing this basic annotation with inference axioms, one can obtain the relations shown with discontinuous arrows: *John* had the intention to *buy*, the *buying* event took place *at the shop* and *John* and the *flowers* were at some point *in the shop*.

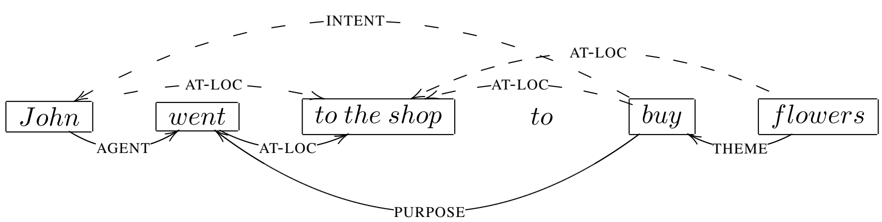

Figure 1: Semantic representation of the sentence *John went to the shop to buy flowers.*

2 Semantic Relations

Semantic relations are the underlying relations between concepts expressed by words or phrases. In other words, semantic relations are implicit associations between concepts in text.

In general, a semantic relation is defined by stating the kind of connection linking two concepts. For example, Hendrickx et al. (2009) loosely define ENTITY-ORIGIN as *an entity is coming or is derived from an origin (e.g., position or material)* and give one example: *Earth is located in the Milky Way*. We find this kind of definition weak and prone to confusion.

Following Helbig (2005), we propose an extended definition for semantic relations, including semantic restrictions for its domain and range. For example, DOMAIN(AGENT) must be an animate concrete object and RANGE(AGENT) must be a situation.

Moreover, we propose to characterize relations by semantic primitives. Primitives indicate if a certain property holds between the arguments of a relation. For example, the primitive *temporal* indicates if the first argument must happen before the second in order for the relation to hold. This primitive holds for CAUSE (a cause must *precede* its effect) and it does not apply to PART-WHOLE since the later relation does not consider time.

Besides having a better understanding of each relation, this extended definition allows us to create a model that automatically obtains inference axioms for composing semantic relations. The model detects possible combinations of relations and identifies the conclusion of composing them.

Formally, we represent a relation R as R(x, y), where R is the relation type and x and y are the first and second argument respectively. R(x, y) should be read *x is* R *of* Y. DOMAIN(R) and RANGE(R) are the sorts of concepts that can be part of the first and second argument respectively. Any ontology can be used to define domains and ranges, e.g., Helbig (2005) defined one to define a set of 89 relations. Primitives are represented by an array P_{R} of length n, where n is the number of primitives and P_{R}^i indicates the value R takes for the ith primitive.

The inverse of R is denoted R^{-1} and can be obtained by simply switching the arguments of R. Given R(x, y), $\mathrm{R}^{-1}(y, x)$ always holds. We can easily define R^{-1} given the definition for R: $\mathrm{DOMAIN}(\mathrm{R}^{-1}) = \mathrm{RANGE}(\mathrm{R})$, $\mathrm{RANGE}(\mathrm{R}^{-1}) = \mathrm{DOMAIN}(\mathrm{R})$, and $P_{\mathrm{R}^{-1}}$ is defined according to the fourth column of Table 1 for each primitive, i.e., $\forall i \in [1, n] : P_{R^{-1}}^i = Inverse(P_R^i)$.

2.1 Semantic Primitives

Relation primitives capture deep characteristics of relations. Huhns and Stephens (1989) define them as:

> They [primitives] are independently determinable for each relation and relatively self-explanatory. They specify a relationship between an element of the domain and an element of the range of the semantic relation being described.

Relation primitives are fundamental properties that cannot be explained using other primitives; they are elemental. They specify basic attributes of a relation by stating if a particular property must hold by definition between the domain and range.

Each relation takes a value for each primitive from the set $V = \{+, -, 0\}$, where '+' indicates that the property holds, '$-$' that it does not hold and '0' that it does not apply. For example, the primitive *volitional* indicates if a relation requires volition between domain and range. AGENT takes as value $+$ for this primitive and PART-WHOLE takes 0.

Primitives complement the definition of a relation by stating if a particular property holds between its arguments. They help to understand the inter-relation differences and clustering relations. Primitives can be used as conditions to be fulfilled in order to determine if a potential relation holds. They are general enough to be determined for a relation, not a particular instantiation. In other words, they state properties that hold for all instances of a relation by definition.

Our set of primitives (Table 1) is inspired on previous work in Knowledge Bases (Huhns and Stephens, 1989). We only select from them useful primitives for our purpose and add more primitives. The additional primitives are justified by the fact that we aim at combining relations capturing semantics

No.	Primitive	Description	Inverse	Ref.
1	Composable	Relation can be meaningfully composed with other relations due to their fundamental characteristics	same	[3]
2	Functional	Domain is in a specific spatial or temporal position with respect to the range in order for the connection to exist	same	[1]
3	Separable	Domain can be temporally or spatially separated from the range, and can thus exist independently of the range	same	[1]
4	Temporal	Domain temporally precedes the range	opposite	[2]
5	Connected	Domain is physically or temporally connected to the range; connection might be indirect.	same	[3]
6	Intrinsic	Relation is an attribute of the essence/stufflike nature of the domain or range	same	[3]
7	Volitional	Relation requires volition between the arguments	same	-
8	Fully Implicational	The existence of the domain implies the existence of the range	opposite	-
9	Weakly Implicational	The existence of the domain generally implies the existence of the range	opposite	-

Table 1: Primitives for characterizing semantic relations, values for the inverse relation and references. In the fifth column, [1] stands for Winston et al. (1987), [2] for Cohen and Losielle (1988) and [3] for Huhns and Stephens (1989). '-' indicates new primitive.

1: **Composable**

R_1 \ R_2	-	0	+
-	×	0	×
0	0	0	0
+	×	0	+

2: **Functional**

R_1 \ R_2	-	0	+
-	-	0	+
0	0	0	0
+	+	0	+

3: **Separable**

R_1 \ R_2	-	0	+
-	-	-	-
0	-	0	+
+	-	+	+

4: **Temporal**

R_1 \ R_2	-	0	+
-	-	-	×
0	-	0	+
+	×	+	+

5: **Connected**

R_1 \ R_2	-	0	+
-	-	-	+
0	-	0	+
+	+	+	+

6: **Intrinsic**

R_1 \ R_2	-	0	+
-	-	0	-
0	0	0	0
+	-	0	+

7: **Volitional**

R_1 \ R_2	-	0	+
-	-	0	+
0	0	0	0
+	+	0	+

8: **F Impl.**

R_1 \ R_2	-	0	+
-	-	0	-
0	0	0	0
+	-	0	+

9: **W Impl.**

R_1 \ R_2	-	0	+
-	-	0	-
0	0	0	0
+	-	0	+

Table 2: Algebra for composing semantic primitives. Each cell of the ith table indicates $P^i_{R_1} \circ P^i_{R_2}$.

from natural language. Whatever the set of chosen relations, it will describe the characteristics of events (who/when/where/how something happened), which elements were involved, connections between events (e.g. CAUSE, CORRELATION). Time (whether an argument is guaranteed to happen before than the other), space and volition (whether or not there must be volition between the arguments) also play an important role.

The fourth column in Table 1 indicates the value of the primitive for the inverse relation. *Same* means the inverse relation takes the same value, *opposite* means it takes the opposite. The opposite of $-$ is $+$, the opposite of $-$ is $+$, and the opposite of 0 is 0.

For example, $P_{\text{AGENT}} = \{+, +, +, 0, -, -, +, 0, 0\}$, indicating that $P^5_{\text{AGENT}} = -$ and $P^7_{\text{AGENT}} = +$, i.e., AGENT(x, y) does not require x and y to be *connected* and it requires *volition* between the arguments. Note that ${P_{\text{AGENT}}}^{-1} = P_{\text{AGENT}}$.

2.2 An Algebra for Composing Semantic Relations

The key to automatically obtaining inference axioms is the ability to know beforehand the result of composing semantic primitives using an algebra. This way, one can identify prohibited combinations of relations and determine conclusions for the composition of valid combinations.

Given $P^i_{\text{R}_1}$ and $P^i_{\text{R}_2}$, i.e., the values of R_1 and R_2 for a primitive p_i, we define an algebra that indicates the result of composing them (i.e., $P^i_{\text{R}_1} \circ P^i_{\text{R}_2}$). Composing two primitives can yield three values: $+$, $-$ or 0, indicating if the primitive holds, does not hold or does not apply to the composition of R_1 and R_2. Additionally, the composition can be prohibited, indicated with $\times$. After composing all the primitives for R_1 and R_2, we obtain the primitives values for the composition of R_1 and R_2 (i.e., $P_{\text{R}_1} \circ P_{\text{R}_2}$).

We define the values for the composition using a table for each primitive. Table 2 depicts the whole algebra. The ith table indicates the rules for composing the ith primitive. For example, regarding the *intrinsic* primitive, we have the following rules:

- If both relations are *intrinsic*, the composition is *intrinsic*;
- else if *intrinsic* does not apply to either relation, the primitive does not apply to the composition;
- else, the composition is not *intrinsic*.

Other rules stated by the algebra are: (1) two relations shall not compose if they have different opposite values for the primitive *temporal*; (2) the composition of R_1 and R_2 is not *separable* if either relation is not *separable*; and (3) if either R_1 or R_2 are *connected*, then the composition is *connected*.

3 Necessary Conditions for Composing Semantic Relations

In principle, one could define axioms for every single possible combination of relations. However, there are two necessary conditions in order to compose R_1 and R_2:

1. They have to be compatible. A pair of relations is compatible if it is possible, from a theoretical point of view, to compose them. Formally, R_1 and R_2 are compatible iff $\text{RANGE}(\text{R}_1) \cap \text{DOMAIN}(\text{R}_2) \neq \emptyset$.
2. A third relation R_3 must fit as conclusion, that is, $\exists \text{R}_3$ such that $\text{DOMAIN}(\text{R}_3) \cap \text{DOMAIN}(\text{R}_1) \neq \emptyset$ and $\text{RANGE}(\text{R}_3) \cap \text{RANGE}(\text{R}_2) \neq \emptyset$.
   Furthermore, P_{R_3} must be compatible with the result of composing P_{R_1} and P_{R_2}.

It is important to note that domain and range compatibility is not enough to compose two relations. For example, given KINSHIP(*Mary*, *John*) and AT-LOCATION(*John*, *Dallas*), no relation can be inferred between *Mary* and *Dallas*.

4 Inference Axioms

An axiom is defined as a set of relations called premises and a conclusion. The composition operator $\circ$ is the basic way of combining two relations to form an axiom. We denote an inference axiom as $\text{R}_1(x, y) \circ \text{R}_2(y, z) \rightarrow \text{R}_3(x, z)$, where R_1 and R_2 are the premises and R_3 the conclusion. In order to instantiate an axiom the premises must have an argument in common, y.

In general, for n relations there are $\binom{n}{2} = \frac{n(n-1)}{2}$ different pairs. For each pair, taking into account the two relations and their inverses, there are $4 \times 4 = 16$ different possible combinations.

We note that $\text{R}_1 \circ \text{R}_2 = (\text{R}_2{}^{-1} \circ \text{R}_1{}^{-1})^{-1}$, reducing the total number of different combinations to 10. Out of these 10, (1) 4 combine R_1, R_2 and their inverses (Table 3); (2) 3 combine R_1 and its inverse; and (3) 3 combine R_2 and its inverse. The most interesting combinations to use as premises for an axiom fall into category (1), since the other two can be resolved by the transitivity property of a relation and its inverse. Therefore, for n relations there are $2n^2 + n$ potential axioms: $\binom{n}{2} \times 4 + 3n = 2 \times n(n-1) + 3n = 2n^2 - 2n + 3n = 2n^2 + n$.

4.1 An Algorithm for Obtaining Inference Axioms

Given a set of relations R defined using the extended definition, one can automatically obtain inference axioms using the following steps for each pair of relations $\text{R}_1 \in R$ and $\text{R}_2 \in R$, where $\text{R}_1 \neq \text{R}_2$:

$R_1 \circ R_2$	$R_1^{-1} \circ R_2$	$R_2 \circ R_1$	$R_2 \circ R_1^{-1}$
$\begin{tikzcd} x \arrow[r, "R_1"] \arrow[dr, "R_3"'] & y \arrow[d, "R_2"] \\ & z \end{tikzcd}$	$\begin{tikzcd} x \arrow[dr, "R_3"'] & y \arrow[l, "R_1"'] \arrow[d, "R_2"] \\ & z \end{tikzcd}$	$\begin{tikzcd} x \arrow[d, "R_2"'] \arrow[dr, "R_3"] & \\ y \arrow[r, "R_1"'] & z \end{tikzcd}$	$\begin{tikzcd} x \arrow[d, "R_2"'] \arrow[dr, "R_3"] & \\ y & z \arrow[l, "R_1"] \end{tikzcd}$

Table 3: The four unique axioms taking as premises R_1 and R_2. R_3 indicates the conclusion.

Role	**Primitive**									**Role**	**Primitive**								
	Composable	Functional	Separable	Temporal	Connected	Intrinsic	Volitional	Fully Impl.	Weakly Impl.		Composable	Functional	Separable	Temporal	Connected	Intrinsic	Volitional	Fully Impl.	Weakly Impl.
ARG0	+	+	+	0	-	-	+	0	0	ARG0^{-1}	+	+	+	0	-	-	+	0	0
ARG1	+	-	+	0	-	-	-	0	0	ARG1^{-1}	+	-	+	0	-	-	-	0	0
MLOC	+	+	0	0	+	-	0	0	0	MLOC^{-1}	+	+	0	0	+	-	0	0	0
MCAU	+	+	+	+	-	+	0	+	+	MCAU^{-1}	+	+	+	-	-	+	0	-	-
MTMP	+	+	0	0	+	-	0	0	0	MTMP^{-1}	+	+	0	0	+	-	0	0	0
MPNC	+	-	+	-	-	-	-	0	-	MPNC^{-1}	+	-	+	+	-	-	-	0	+
MMNR	+	-	+	0	-	-	+	0	0	MMNR^{-1}	+	-	+	0	-	-	+	0	0

Table 4: Semantic Roles in PropBank, their inverses and their primitives.

Repeat Steps 1, 2 and 3 for $(R_i, R_j) \in [(R_1, R_2), (R_1^{-1}, R_2), (R_2, R_1), (R_2, R_1^{-1})]$:

1. **Domain and range compatibility**
   If $\text{RANGE}(R_i) \cap \text{DOMAIN}(R_j) = \emptyset$, **break**
2. **Primitives composition**
   Using the algebra for composing semantic primitives, calculate $P_{R_i} \circ P_{R_j}$
3. **Conclusion match** Repeat for $R_3 \in R$
   If $\text{DOMAIN}(R_3) \cap \text{DOMAIN}(R_i) \neq \emptyset$ **and** $\text{RANGE}(R_3) \cap \text{RANGE}(R_j) \neq \emptyset$
   and $consistent(P_{R_3}, P_{R_i} \circ P_{R_j})$, **then**
   $inference\_axioms$ += $R_i(x, y) \circ R_j(y, z) \rightarrow R_3(x, z)$

The method $consistent(P_1, P_2)$ is a simple procedure that compares the values assigned to each primitive one by one. Two values for the same primitive are compatible unless they have different opposites or either value is '$\times$' (i.e., prohibited).

5 Case Study: PropBank

PropBank (Palmer et al., 2005) adds a layer of predicate-argument information, or semantic role labels, on top of the syntactic trees provided by the Penn TreeBank. Along with FrameNet, it is the resource most widely used for semantic role annotation.

PropBank uses a series of numeric core roles (ARG0 - ARG5) and a set of more general roles, ARGMs (e.g. MTMP, MLOC, MMNR). The interpretation of the numeric roles is determined by a verb-specific framesets, although ARG0 and ARG1 usually correspond to the prototypical AGENT and THEME. On the other hand, the meaning of AGRMs generalize across verbs.

An example of PropBank annotation is the following: [Winston]$_{\text{ARG0}}$ [procrastinated]$_{rel}$ [a lot]$_{\text{MADV}}$ [due to his nervous demeanor]$_{\text{MCAU}}$. Palmer et al. (2005) discuss the creation of PropBank. For more information about the semantics of each role, we refer the reader to the annotation guidelines[1].

Since ARG2, AGR3, ARG4 and ARG5 do not have a common meaning across verbs, they become not *composable*. For example, ARG2 is used for INSTRUMENT in the frameset *kick.01* and for BENEFACTIVE in the frameset *call.02*.

[1]http://verbs.colorado.edu/˜m palmer/projects/ace/PBguidelines.pdf

$\mathbf{R}_1$	$\mathbf{R}_2$: a: ARG0^{-1}	b: ARG1^{-1}	c: MLOC^{-1}	d: MCAU^{-1}	e: MTMP^{-1}	f: MPNC^{-1}	g: MMNR^{-1}
a: ARG0	=	-	-	a	-	a	-
b: ARG1	-	=	-	-	-	b	-
c: MLOC	-	-	=	c	-	c	-
d: MCAU	a	-	c	=	e	-	-
e: MTMP	-	-	-	e	=	e	-
f: MPNC	a	b	c	-	e	=	g
g: MMNR	-	-	-	-	-	g	=

Table 5: Results after applying the steps depicted in Section 4.1 using PropBank semantic roles. A letter indicates an inference axiom $\text{R}_1 \circ \text{R}_2 \rightarrow \text{R}_3$ by indicating the conclusion R_3. '-' indicates that the combination is not prohibited but a relation compatible with $P_{\text{R}_1} \circ P_{\text{R}_2}$ could not be found; '=' indicates that the cell corresponds to a relation and its inverse.

The remaining labels (ARG0, ARG1 and all ARGMs) do generalize in meaning across verbs. Roles MEXT, MDIS, MADV, MNEG, MMOD, MDIR, are not *composable* because they encode a very narrow semantic connection. Manual examination of several examples leads to this conclusion.

Table 4 depicts the primitives for the roles which are *composable* and their inverses. Note that for any two relations their primitives are different.

PropBank does not provide domains and ranges for its roles, although we can specify our own. We do so by using the ontology defined by Helbig (2005). All relations in PropBank are denoted as $\text{R}(x, y)$, where x is an argument of y, and y is a verb. The range of all relations is a situation. The domain of AGR0 and ARG1 are objects, the domain of MLOC and MTMP local and temporal descriptors respectively, the domain of MMNR qualities or states, and the domain of MPNC and MCAU are situations.

5.1 Inference Axioms from PropBank

Out of the four possible axioms between any pair of relations (Table 3), the only way to compose two relations from PropBank is by using as common argument y a verb. This restriction is due to the fact that PropBank exclusively annotates relations between a verb and its arguments. Thus, the only possible axiom for any pair of roles R_1 and R_2 is $\text{R}_1(x, y) \circ \text{R}_2{}^{-1}(y, z) \rightarrow \text{R}_3(x, z)$, where y is a verb.

Table 5 shows the eight inference axioms obtained after following the steps depicted in Section 4.1. Note that the matrix is symmetric as stated by the property $\text{R}_1 \circ \text{R}_2 = (\text{R}_2{}^{-1} \circ \text{R}_1{}^{-1})^{-1}$.

Some of the axioms obtained are:

- MCAU $\circ$ MLOC^{-1} $\rightarrow$ MLOC^{-1}, *the location of a cause is the same than the location of its effect.*
- MPNC $\circ$ ARG0^{-1} $\rightarrow$ ARG0^{-1}, *the agent of an action is inherited by its purpose.*
- MPNC $\circ$ MMNR^{-1} $\rightarrow$ MMNR^{-1}, *the manner of an action is inherited by its purpose.*

5.2 Evaluation

First, we evaluated all the instantiations of axiom MPNC $\circ$ MMNR^{-1} $\rightarrow$ MMNR^{-1}. This axiom can be instantiated 237 times using PropBank annotation, yielding 189 new MANNER not present in PropBank. The overall accuracy is 0.797, superior to state-of-the art semantic role labelers.

Second, we have evaluated the accuracy of the eight inference axioms (Table 5). Since PropBank is a large corpus, the amount of instantiations found for all axioms is too large to be checked by hand. We have manually evaluated the first 1,000 sentences that are an instantiation of any axiom. Since a sentence may instantiate several axioms, we have actually evaluated 1,412 instantiations. The first 1,000 sentences which are an instantiation of any axiom are found within the first 31,450 sentences in PropBank. Table 6 shows the number of roles PropBank annotates for these sentences.

Role	No. Instances
CAUSE	421
PURPOSE	768
AGENT	22,525
THEME	29,738
AT-LOCATION	2,024
AT-TIME	5,743
MANNER	2,212

Table 6: Number of relations in PropBank for the first 31,450 sentences.

		no heuristic			with heuristic		
No.	Axiom	No. Inst.	Acc.	Produc.	No. Inst.	Acc.	Produc.
1	CAU $\circ$ AGT^{-1} $\rightarrow$ AGT^{-1}	201	0.40	0.89%	75	0.67	0.33%
2	CAU $\circ$ AT-L $\rightarrow$ AT-L	17	0.82	0.84%	15	0.93	0.74%
3	CAU $\circ$ AT-T $\rightarrow$ AT-T	72	0.85	1.25%	69	0.87	1.20%
1-3	CAU $\circ$ R$_2$ $\rightarrow$ R$_3$	290	0.53	0.96%	159	0.78	0.53%
4	PRP $\circ$ AGT^{-1} $\rightarrow$ AGT^{-1}	375	0.89	1.66%	347	0.94	1.54%
5	PRP $\circ$ THM^{-1} $\rightarrow$ THM^{-1}	489	0.12	1.64%	87	0.65	0.29%
6	PRP $\circ$ AT-L $\rightarrow$ AT-L	49	0.90	2.42%	48	0.92	2.37%
7	PRP $\circ$ AT-T $\rightarrow$ AT-T	138	0.84	2.40%	129	0.88	2.25%
8	PRP $\circ$ MNR^{-1} $\rightarrow$ MNR^{-1}	71	0.82	3.21%	70	0.83	3.16%
4-8	PRP $\circ$ R$_2$ $\rightarrow$ R$_3$	1,122	0.54	1.80%	681	0.88	1.09%
1-8	**All**	1,412	0.54	2.26%	**840**	**0.86**	**1.35%**

Table 7: Axioms used during evaluation, number of instances, accuracy and productivity. Results are reported both using and not using the heuristic. Productivity refers to the number of relations added by the axiom in relative terms.

Table 7 depicts the total number of instantiations for each axiom and its accuracy (columns 3 and 4). Accuracies range from 0.12 to 0.90, showing that the plausibility of an axiom depends on the axiom. The average accuracy for axioms involving MCAU is 0.53 and for axioms involving MPNC is 0.54.

Axiom MCAU $\circ$ ARG0^{-1} $\rightarrow$ ARG0^{-1} adds 201 relations, which corresponds to 0.89% in relative terms. Its accuracy is low, 0.40. Other axioms are less productive overall, but have a greater relative impact and accuracy. For example, axiom MPNC $\circ$ MMNR^{-1} $\rightarrow$ MMNR^{-1}, only yields 71 new MMNR, and yet it is adding 3.21% in relative terms with an accuracy of 0.82.

It is worth noting that overall, applying the eight axioms used during evaluation adds 1,412 relations on top of the ones already present (2.26% in relative terms) with an accuracy of 0.54.

5.3 Error Analysis

Because of the low accuracy of axioms 1 and 5, an error analysis was performed. We found that unlike other axioms, these axioms often yield a relation type that is already present in the semantic representation. Specifically, axioms 1 and 5 often yield R(x, z) when R(x', z) is already known.

An example can be found in Figure 4, where axiom 5 yields ARG1(*orders*, *to buy*) when the relation ARG1(*the basket*, *to buy*) is already present. We use the following heuristic in order to improve the accuracy of axioms 1 and 5: *do not instantiate an axiom* R$_1$*(x, y)* $\circ$ R$_2$*(y, z)* $\rightarrow$ R$_3$*(x, z) if a relation of the form* R$_3$*(x', z) is already known.*

This simple heuristic allows us to augment the accuracy of the inferences at the cost of lowering their productivity. The last three columns in Table 7 show results when using the heuristic. The eight axioms add 840 relations (1.35% in relative terms) with an accuracy of 0.86.

5.4 Examples

In this section we present several examples of instantiations. We provide the full text of each example, but only the relevant semantic annotation for instantiating axioms. For all examples, solid arrows indicate semantic role annotation from PropBank, and discontinuous arrows inferred relations.

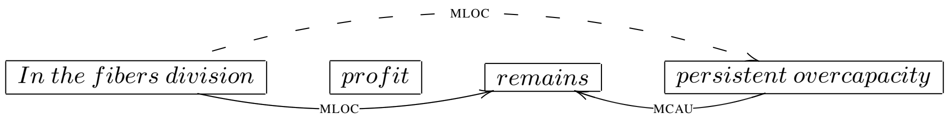

Figure 2: *In the fibers division, profit remains weak, largely because of persistent overcapacity.* (wsj_0552, 28).

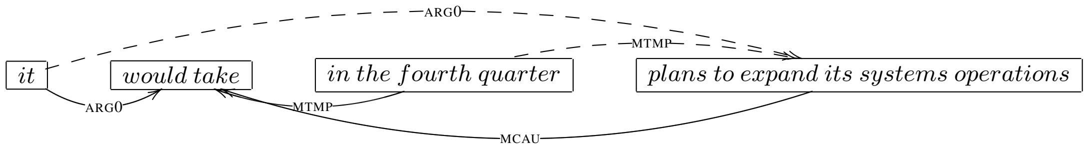

Figure 3: *First Tennessee National Corp. said it would take a $4 million charge in the fourth quarter, as a result of plans to expand its systems operation.* (wsj_0621, 0).

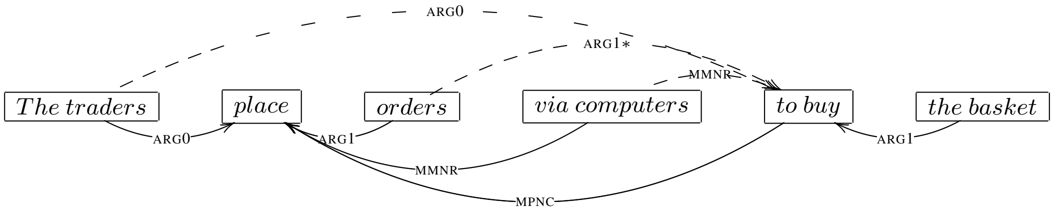

Figure 4: *When it occurs, the traders place orders via computers to buy the basket of stocks . . . in whichever market is cheaper and sell them in the more expensive market; . . .* (wsj_0118, 48).

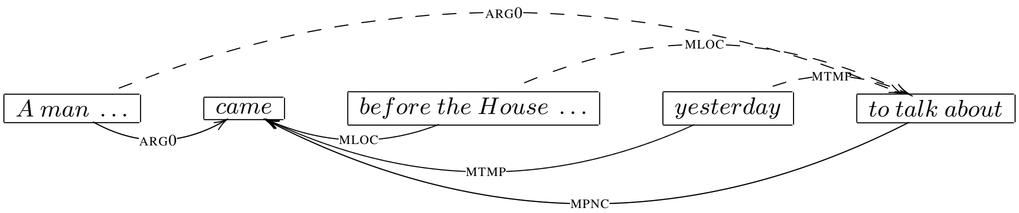

Figure 5: *A man from the Bush administration came before the House Agriculture Committee yesterday to talk about . . .* (wsj_0134, 0).

Figures 2 and 3 instantiate axioms 1, 2 and 3. For these examples, all inferences are correct.

Figures 4 and 5 instantiate the rest of axioms. Not using the heuristic leads to a wrong inference in the example shown in Figure 4, indicated with *. Using the heuristic, all inferences are correct.

6 Comparison with Previous Work

There have been abundant proposals to detect semantic relations without taking into account composition of relations. All these approaches, regardless of their particular details, take as their input text and output the relations found in it. In contrast, the framework proposed in this article obtains axioms that take as their input relations found in text and output more relations previously ignored.

Generally, efforts to extract semantic relations have concentrated on particular sets of relations or a single relation, e.g. CAUSE (Bethard and Martin, 2008; Chang and Choi, 2006) and PART-WHOLE (Girju et al., 2006). Automatic detection of semantic roles has received a lot of attention lately (Màrquez et al., 2008; Carreras and Màrquez, 2005). The SemEval-2007 Task 04 (Girju et al., 2007) and SemEval-2010 Task 08 (Hendrickx et al., 2009) aimed at relations between nominals. There has been work on detecting relations within noun phrases (Moldovan et al., 2004; Nulty, 2007), clauses (Szpakowicz et al., 1995) and syntax-based comma resolution (Srikumar et al., 2008).

Previous research has exploited the idea of using semantic primitives to define and classify semantic relations under different names. Among others, the literature uses *relation elements*, *deep structure*, *aspects* and *primitives*. To the best of our knowledge, the first effort on describing semantic relations using primitives was made by Chaffin and Herrmann (1987). They introduce Relation Element Theory, and differentiate relations by *relation elements*. The authors describe a set of 31 relations clustered in five groups (CONTRAST, SIMILARS, CLASS INCLUSION, CASE-RELATIONS, PART-WHOLE), and distinguish each relation by its *relations elements* and not just a definition and examples. Their 30 *relation elements* are clustered into five groups (elements of intensional force, dimension elements, elements of agreement, propositional elements, elements of part-whole inclusion). They only use the *elements* to define relations, not to compose relations.

Winston et al. (1987) work with six subtypes of PART-WHOLE and uses 3 relation elements (*functional*, *homeomerous* and *separable*) to distinguish the subtypes. Cohen and Losielle (1988) introduce the notion of *deep structure* and characterize it using two aspects: *hierarchical* and *temporal*. Huhns and Stephens (1989) extend previous works by considering an extended set of 10 primitives.

In Computational Linguistics there have been previous proposals to combine semantic relations. Harabagiu and Moldovan (1998) manually extract plausible inference axioms using WordNet relations. Helbig (2005) transforms chains of relations into theoretical axioms. On the other hand, the model presented in this paper extracts inference axioms automatically.

Composing relations has been proposed before in the more general field of Artificial Intelligence, in particular in the context of Knowledge Bases. Cohen and Losielle (1988) point out that two relations shall combine if and only if they do not have contradictory values for the aspect *hierarchical* or *temporal*. They work with a set of nine specific relations (CAUSES, COMPONENT-OF, FOCUS-OF, MECHANISM-OF, PRODUCT-OF, PURPOSE-OF, SETTING-OF, SUBJECT-OF and SUBFIELD-OF) and their inverses. Huhns and Stephens (1989) are the first to propose an algebra for composing semantic primitives. Unlike ours, their set of relations is not linguistically motivated; ten of them map to some sort of PART-WHOLE (e.g. PIECE-OF, SUBREGION-OF).

7 Conclusions

In this paper, we have presented a model to compose semantic relations. The model is independent of any particular set of relations and is able to obtain inference axioms. These axioms take as their input two semantic relations and yield a previously ignored relation as conclusion.

The model is based on an extended definition of semantic relations, including restrictions on domains and ranges and values for a set of semantic primitives. We have defined an algebra for composing semantic primitives. This algebra is the key to automatically identify the resulting relation of composing a pair of compatible relations and to form an axiom.

The proposed algorithm to compose semantic relations identifies eight inference axioms using PropBank relations. When instantiated in a subset of PropBank, these axioms add 2.26% of annotation in relative terms with an accuracy of 0.54. We believe these results are worthwhile for a completely unsupervised approach to obtain semantic relations. Adding a simple heuristic improves the accuracy to 0.86, lowering the productivity in relative terms to 1.35%.

The model has limitations and is not always correct. First, relations are defined manually and mistakes could be made when assigning values to their primitives. Second, the algebra for composing primitives is also manually defined.

We find the first problem easy to overcome. Whatever the set of relations one might use, we believe thinking in terms of primitives helps to understand the nature of the relations and their differences. An issue might be that the proposed set of primitives is not enough for a particular set, but more primitives could be added to solve this eventuality.

A further issue with the algebra is the fact that primitives are composed orthogonally. This is a simplification, but we have shown that this simplified algebra works.

Even though different sets of semantic relations may call for different ontologies to define domains and ranges, and possibly an extended set of primitives, we believe the model presented in this paper is applicable to any set. As far as we are concerned, this is a novel way to compose semantic relations in the field of Computational Linguistics.

References

Bethard, S. and J. H. Martin (2008). Learning Semantic Links from a Corpus of Parallel Temporal and Causal Relations. In *Proceedings of ACL-08: HLT, Short Papers*, Columbus, OH.

Carreras, X. and L. Màrquez (2005). Introduction to the CoNLL-2005 shared task: semantic role labeling. In *Proc. of the 9th Conf. on Computational Natural Language Learning*, Morristown, NJ.

Chaffin, R. and D. J. Herrmann (1987). Relation Element Theory: A New Account of the Representation and Processing of Semantic Relations. In D. S. Gorfein and R. R. Hoffman (Eds.), *Memory and Learning. The Ebbinghaus Centennial Conference*.

Chang, D. S. and K. S. Choi (2006). Incremental cue phrase learning and bootstrapping method for causality extraction using cue phrase and word pair probabilities. *Information Processing & Management 42*(3), 662–678.

Cohen, P. R. and C. L. Losielle (1988). Beyond ISA: Structures for Plausible Inference in Semantic Networks. In *Proceedings of the Seventh National conference on Artificial Intelligence*, St. Paul, MN.

Girju, R., A. Badulescu, and D. Moldovan (2006). Automatic Discovery of Part-Whole Relations. *Computational Linguistics 32*(1), 83–135.

Girju, R., P. Nakov, V. Nastase, S. Szpakowicz, P. Turney, and D. Yuret (2007). SemEval-2007 Task 04: Classification of Semantic Relations between Nominals. In *Proceedings of the Fourth International Workshop on Semantic Evaluations*, Prague, Czech Republic.

Harabagiu, S. and D. Moldovan (1998). Knowledge Processing on an Extended WordNet. In C. Fellbaum (Ed.), *WordNet: An Electronic Lexical Database and Some of its Applications*, Chapter 17, pp. 684–714. The MIT Press.

Helbig, H. (2005). *Knowledge Representation and the Semantics of Natural Language* (1st ed.). Springer.

Hendrickx, I., S. N. Kim, Z. Kozareva, P. Nakov, Diarmuid, S. Padó, M. Pennacchiotti, L. Romano, and S. Szpakowicz (2009). SemEval-2010 Task 8: Multi-Way Classification of Semantic Relations Between Pairs of Nominals. In *Proceedings of the Workshop on Semantic Evaluations*, Boulder, CO.

Huhns, M. N. and L. M. Stephens (1989). Plausible Inferencing Using Extended Composition. In *IJCAI'89: Proceedings of the 11th international joint conference on AI*, San Francisco, CA.

Màrquez, L., X. Carreras, K. C. Litkowski, and S. Stevenson (2008, June). Semantic Role Labeling: An Introduction to the Special Issue. *Computational Linguistics 34*(2), 145–159.

Moldovan, D., A. Badulescu, M. Tatu, D. Antohe, and R. Girju (2004). Models for the Semantic Classification of Noun Phrases. In *HLT-NAACL 2004: Workshop on Computational Lexical Semantics*.

Nulty, P. (2007). Semantic Classification of Noun Phrases Using Web Counts and Learning Algorithms. In *Proceedings of the ACL 2007 Student Research Workshop*, Prague, Czech Republic.

Palmer, M., D. Gildea, and P. Kingsbury (2005). The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics 31*(1), 71–106.

Srikumar, V., R. Reichart, M. Sammons, A. Rappoport, and D. Roth (2008). Extraction of Entailed Semantic Relations Through Syntax-Based Comma Resolution. In *Proceedings of ACL-08: HLT*, Columbus, OH.

Szpakowicz, B., K. Barker, and S. Szpakowicz (1995). Interactive semantic analysis of Clause-Level Relationships. In *Proc. of the 2nd Conference of the Pacific Association for Computational Linguistics*.

Winston, M. E., R. Chaffin, and D. Herrmann (1987). A Taxonomy of Part-Whole Relations. *Cognitive Science 11*(4), 417–444.