

Unsupervised Learning of Semantic Relation Composition

Eduardo Blanco and **Dan Moldovan**
Human Language Technology Research Institute
The University of Texas at Dallas
Richardson, TX 75080 USA
{eduardo,moldovan}@hlt.utdallas.edu

Abstract

This paper presents an unsupervised method for deriving inference axioms by composing semantic relations. The method is independent of any particular relation inventory. It relies on describing semantic relations using primitives and manipulating these primitives according to an algebra. The method was tested using a set of eight semantic relations yielding 78 inference axioms which were evaluated over PropBank.

1 Introduction

Capturing the meaning of text is a long term goal within the NLP community. Whereas during the last decade the field has seen syntactic parsers mature and achieve high performance, the progress in semantics has been more modest. Previous research has mostly focused on relations between particular kind of arguments, e.g., semantic roles, noun compounds. Notwithstanding their significance, they target a fairly narrow text semantics compared to the broad semantics encoded in text.

Consider the sentence in Figure 1. Semantic role labelers exclusively detect the relations indicated with solid arrows, which correspond to the sentence syntactic dependencies. On top of those roles, there are at least three more relations (discontinuous arrows) that encode semantics other than the verb-argument relations.

In this paper, we venture beyond semantic relation extraction from text and investigate techniques to compose them. We explore the idea of inferring

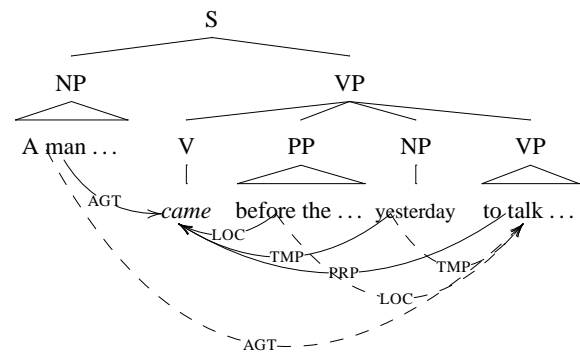


Figure 1: Semantic representation of *A man from the Bush administration came before the House Agricultural Committee yesterday to talk about ...* (wsj_0134, 0).

a new relation linking the ends of a chain of relations. This scheme, informally used previously for combining HYPERNYM with other relations, has not been studied for arbitrary pairs of relations.

For example, it seems adequate to state the following: if x is PART-OF y and y is HYPERNYM of z , then x is PART-OF z . An inference using this rule can be obtained instantiating x , y and z with *engine*, *car* and *convertible*. Going a step further, we consider nonobvious inferences involving AGENT, PURPOSE and other semantic relations.

The novelties of this paper are twofold. First, an extended definition for semantic relations is proposed, including (1) semantic restrictions for their domains and ranges, and (2) semantic primitives.

Second, an algorithm for obtaining inference axioms is described. Axioms take as their premises chains of two relations and output a new relation linking the ends of the chain. This adds an extra layer of semantics on top of previously extracted re-

Primitive	Description	Inv.	Ref.
1: Composable	Relation can be meaningfully composed with other relations due to their fundamental characteristics	id.	[3]
2: Functional	x is in a specific spatial or temporal position with respect to y in order for the connection to exist	id.	[1]
3: Homeomeric	x must be the same kind of thing as y	id.	[1]
4: Separable	x can be temporally or spatially separated from y ; they can exist independently	id.	[1]
5: Temporal	x temporally precedes y	op.	[2]
6: Connected	x is physically or temporally connected to y ; connection might be indirect.	id.	[3]
7: Intrinsic	Relation is an attribute of the essence/stufflike nature of x and y	id.	[3]
8: Volitional	Relation requires volition between the arguments	id.	-
9: Universal	Relation is always true between x and y	id.	-
10: Fully Implicational	The existence of x implies the existence of y	op.	-
11: Weakly Implicational	The existence of x sometimes implies the existence of y	op.	-

Table 1: List of semantic primitives. In the fourth column, [1] stands for (Winston et al., 1987), [2] for (Cohen and Losielle, 1988) and [3] for (Huhns and Stephens, 1989).

lations. The conclusion of an axiom is identified using an algebra for composing semantic primitives.

We name this framework Composition of Semantic Relations (CSR). The extended definition, set of primitives, algebra to compose primitives and CSR algorithm are independent of any particular set of relations. We first presented CSR and used it over PropBank in (Blanco and Moldovan, 2011). In this paper, we extend that work using a different set of primitives and relations. Seventy eight inference axioms are obtained and an empirical evaluation shows that inferred relations have high accuracies.

2 Semantic Relations

Semantic relations are underlying relations between concepts. In general, they are defined by a textual definition accompanied by a few examples. For example, Chklovski and Pantel (2004) loosely define ENABLEMENT as a relation that holds between two verbs V_1 and V_2 when the pair can be glossed as V_1 is accomplished by V_2 and gives two examples: *assess::review* and *accomplish::complete*.

We find this widespread kind of definition weak and prone to confusion. Following (Helbig, 2005), we propose an extended definition for semantic relations, including semantic restrictions for its arguments. For example, AGENT(x, y) holds between an animate concrete object x and a situation y .

Moreover, we propose to characterize relations by *semantic primitives*. Primitives indicate whether a property holds between the arguments of a relation,

e.g., the primitive *temporal* indicates if the first argument must happen before the second.

Besides having a better understanding of each relation, this extended definition allows us to identify possible and not possible combinations of relations, as well as to automatically determine the conclusion of composing a possible combination.

Formally, for a relation $R(x, y)$, the extended definitions specifies: (a) DOMAIN(R) and RANGE(R) (i.e., semantic restrictions for x and y); and (b) P_R (i.e., values for the primitives). The inverse relation R^{-1} can be obtained by switching domain and range, and defining $P_{R^{-1}}$ as depicted in Table 1.

2.1 Semantic Primitives

Semantic primitives capture deep characteristics of relations. They are independently determinable for each relation and specify a property between an element of the domain and an element of the range of the relation being described (Huhns and Stephens, 1989). Primitives are fundamental, they cannot be explained using other primitives.

For each primitive, each relation takes a value from the set $V = \{+, -, 0\}$. ‘+’ indicates that the primitive holds, ‘-’ that it does not hold, and ‘0’ that it does not apply. Since a cause must precede its effect, we have $P_{\text{CAUSE}}^{\text{temporal}} = +$.

Primitives complement the definition of a relation and completely characterize it. Coupled with domain and range restrictions, primitives allow us to automatically manipulate and reason over relations.

1:Composable			2:Functional			3:Homeomorous			4:Separable			5:Temporal			6:Connected				
R ₁	R ₂		R ₁	R ₂		R ₁	R ₂		R ₁	R ₂		R ₁	R ₂		R ₁	R ₂			
	-	0		+	-		0	+		-	0		+	-		0	+	-	0
-	×	0	×	-	-	0	+	-	-	-	-	-	-	×	-	-	+		
0	0	0	0	0	0	0	0	0	-	0	+	0	-	0	+	0	-	0	+
+	×	0	+	+	0	+	+	+	-	0	+	+	×	+	+	+	+	+	+

7:Intrinsic			8:Volitional			9:Universal			10:F. Impl.			11:W. Impl.					
R ₁	R ₂		R ₁	R ₂		R ₁	R ₂		R ₁	R ₂		R ₁	R ₂				
	-	0		+	-		0	+		-	0		+	-	0	+	
-	-	0	-	-	0	-	-	0	-	-	0	×	-	-	-	×	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	-	0	+	
+	-	0	+	+	+	+	-	0	+	+	×	0	+	+	×	+	+

Table 2: Algebra for composing semantic primitives.

The set of primitives used in this paper (Table 1) is heavily based on previous work in Knowledge Bases (Huhns and Stephens, 1989), but we considered some new primitives. The new primitives are justified by the fact that we aim at composing relations capturing the semantics from natural language. Whatever the set of relations, it will describe the characteristics of events (who / what / where / when / why / how) and connections between them (e.g., CAUSE, CORRELATION). Time, space and volition also play an important role. The third column in Table 1 indicates the value of the primitive for the inverse relation: *id.* means it takes the same; *op.* the opposite. The opposite of $-$ is $+$, the opposite of $+$ is $-$, and the opposite of 0 is 0.

2.1.1 An Algebra for Composing Semantic Primitives

The key to automatically obtain inference axioms is the ability to know the result of composing primitives. Given $P_{R_1}^i$ and $P_{R_2}^i$, i.e., the values of the i th primitive for R_1 and R_2 , we define an algebra for $P_{R_1}^i \circ P_{R_2}^i$, i.e., the result of composing them. Table 2 depicts the algebra for all primitives. An ‘ \times ’ means that the composition is prohibited.

Consider, for example, the Intrinsic primitive: if both relations are *intrinsic* ($+$), the composition is *intrinsic* ($+$); else if *intrinsic* does not apply to either relation (0), the primitive does not apply to the composition either (0); else the composition is not *intrinsic* ($-$).

3 Inference Axioms

Semantic relations are composed using inference axioms. An axiom is defined by using the composi-




Table 3: The four unique possible axioms taking as premises R_1 and R_2 . Conclusions are indicated by R_3 and are not guaranteed to be the same for the four axioms.

tion operator ‘ \circ ’; it combines two relations called *premises* and yields a *conclusion*. We denote an axiom as $R_1(x, y) \circ R_2(y, z) \rightarrow R_3(x, z)$, where R_1 and R_2 are the premises and R_3 the conclusion. In order to instantiate an axiom, the premises must form a chain by having argument y in common.

In general, for n relations there are $\binom{n}{2}$ pairs. For each pair, taking into account inverse relations, there are 16 possible combinations. Applying property $R_i \circ R_j = (R_j^{-1} \circ R_i^{-1})^{-1}$, only 10 are unique: (a) 4 combine R_1 , R_2 and their inverses (Table 3); (b) 3 combine R_1 and R_1^{-1} ; and (c) 3 combine R_2 and R_2^{-1} . The most interesting axioms fall into category (a) and there are $\binom{n}{2} \times 4 + 3n = 2 \times n(n-1) + 3n = 2n^2 + n$ potential axioms in this category.

Depending on n , the number of potential axioms to consider can be significantly large. For $n = 20$, there are 820 axioms to explore and for $n = 30$, 1,830. Manual examination of those potential ax-

Relation R		Domain	Range	P_R^1	P_R^2	P_R^3	P_R^4	P_R^5	P_R^6	P_R^7	P_R^8	P_R^9	P_R^{10}	P_R^{11}
a: CAU	CAUSE	si	si	+	+	-	+	+	-	+	0	-	+	+
b: INT	INTENT	si	aco	+	+	-	+	-	-	-	+	-	0	-
c: PRP	PURPOSE	si, ao	si, co, ao	+	-	-	+	-	-	-	-	-	0	-
d: AGT	AGENT	aco	si	+	+	-	+	0	-	-	+	-	0	0
e: MNR	MANNER	st, ao, ql	si	+	-	-	+	0	-	-	+	-	0	0
f: AT-L	AT-LOCATION	o, si	loc	+	+	-	0	0	+	-	0	-	0	0
g: AT-T	AT-TIME	o, si	tmp	+	+	-	0	0	+	-	0	-	0	0
h: SYN	SYNONYMY	ent	ent	+	-	+	0	0	0	+	0	+	0	0

Table 4: Extended definition for the set of relations.

ioms would be time-consuming and prone to errors. We avoid this by using the extended definition and the algebra for composing primitives.

3.1 Necessary Conditions for Composing Semantic Relations

There are two necessary conditions for composing R_1 and R_2 :

- 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 $\text{RANGE}(R_1) \cap \text{DOMAIN}(R_2) \neq \emptyset$.
- A third relation R_3 must match as conclusion, i.e., $\exists R_3$ such that $\text{DOMAIN}(R_3) \cap \text{DOMAIN}(R_1) \neq \emptyset$ and $\text{RANGE}(R_3) \cap \text{RANGE}(R_2) \neq \emptyset$. Furthermore, P_{R_3} must be consistent with $P_{R_1} \circ P_{R_2}$.

3.2 CSR: An Algorithm for Composing Semantic Relations

Consider any set of relations R defined using the extended definition. One can obtain inference axioms using the following algorithm:

For $(R_1, R_2) \in R \times R$:

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. **Conclusion match**
Repeat for $R_3 \in \text{possible_conc}(R, R_i, R_j)$:
 - (a) **If** $\text{DOMAIN}(R_3) \cap \text{DOMAIN}(R_i) = \emptyset$ **or** $\text{RANGE}(R_3) \cap \text{RANGE}(R_j) = \emptyset$, **break**
 - (b) **If** $\text{consistent}(P_{R_3}, P_{R_i} \circ P_{R_j})$,
 $\text{axioms} += R_i(x, y) \circ R_j(y, z) \rightarrow R_3(x, z)$

Given R , R^{-1} can be automatically obtained (Section 2). $\text{Possible_conc}(R, R_i, R_j)$ returns the set R

unless $R_i(R_j)$ is universal ($P^9 = +$), in which case it returns $R_j(R_i)$. $\text{Consistent}(P_{R_1}, P_{R_2})$ is a simple procedure that compares the values assigned to each primitive; two values are consistent unless they have different opposite values or any of them is ‘ \times ’ (i.e., the composition is prohibited).

3.3 An Example: Agent and Purpose

We present an example of applying the CSR algorithm by inspecting the potential axiom $\text{AGENT}(x, y) \circ \text{PURPOSE}^{-1}(y, z) \rightarrow R_3(x, z)$, where x is the agent of y , and action y has as its purpose z . A statement instantiating the premises is $[Mary]_x [came]_y$ to $[talk]_z$ about the issue. Knowing $\text{AGENT}(Mary, came)$ and $\text{PURPOSE}^{-1}(came, talk)$, our goal is to identify the links $R_3(Mary, talk)$, if any.

We use the relations as defined in Table 4. First, we note that both AGENT and PURPOSE^{-1} are compatible (Step 1). Second, we must identify the possible conclusions R_3 that fit as conclusions (Step 2).

Given P_{AGENT} and $P_{\text{PURPOSE}^{-1}}$, we obtain $P_{\text{AGENT}} \circ P_{\text{PURPOSE}^{-1}}$ using the algebra:

$$\begin{aligned}
P_{\text{AGENT}} &= \{+, +, -, +, 0, -, -, +, -, 0, 0\} \\
P_{\text{PURPOSE}^{-1}} &= \{+, -, -, +, +, -, -, -, 0, +\} \\
P_{\text{AGENT}} \circ P_{\text{PURPOSE}^{-1}} &= \{+, +, -, +, +, -, -, +, -, 0, +\}
\end{aligned}$$

Out of all relations (Section 4), AGENT and INTENT^{-1} fit the conclusion match. First, their domains and ranges are compatible with the composition (Step 2a). Second, both P_{AGENT} and $P_{\text{INTENT}^{-1}}$ are consistent with $P_{\text{AGENT}} \circ P_{\text{PURPOSE}^{-1}}$ (Step 2b). Thus, we obtain the following axioms: $\text{AGENT}(x, y) \circ \text{PURPOSE}^{-1}(y, z) \rightarrow \text{AGENT}(x, z)$ and $\text{AGENT}(x, y) \circ \text{PURPOSE}^{-1}(y, z) \rightarrow \text{INTENT}^{-1}(x, z)$.

Instantiating the axioms over $[Mary]_x [came]_y$ to $[talk]_z$ about the issue yields $\text{AGENT}(Mary, talk)$ and $\text{INTENT}^{-1}(Mary, talk)$. Namely, the axioms

		R ₂							
R ₁	a	b	c	d	e	f	g	h	
a	a	:	:	-	f	g	a		
b				-	f	g	b		
c	:	b	c	-	e	f	g	c	
d	d	-	d		d	f	g	d	
e	-	b	e		e	f	g	e	
f								f	
g								g	
h	a	b	c	d	e	f	g	h	

		R ₂							
R ₁	a	b	c	d	e	f	g	h	
a ⁻¹	:	b	b		-	f	g	a ⁻¹	
b ⁻¹	b ⁻¹	:	:		b ⁻¹ ,d ⁻¹	f	g	b ⁻¹	
c ⁻¹	b ⁻¹	:	:		e	f	g	c ⁻¹	
d ⁻¹				-		f	g	d ⁻¹	
e ⁻¹	-	b,d	e ⁻¹		e,e ⁻¹	f	g	e ⁻¹	
f ⁻¹	f ⁻¹	f ⁻¹	f ⁻¹	f ⁻¹	f ⁻¹	-	-	f ⁻¹	
g ⁻¹	g ⁻¹	g ⁻¹	g ⁻¹	g ⁻¹	g ⁻¹	-	-	g ⁻¹	
h ⁻¹	a	b	c	d	e	f	g	h,h ⁻¹	

		R ₂							
R ₁	a ⁻¹	b ⁻¹	c ⁻¹	d ⁻¹	e ⁻¹	f ⁻¹	g ⁻¹	h ⁻¹	
a	:		:	d ⁻¹	-			a	
b		:	:					b	
c	:	:	:	b,d ⁻¹	e ⁻¹			c	
d	d		b ⁻¹ ,d	-	b,d			d	
e	-		e	b ⁻¹ ,d ⁻¹	e,e ⁻¹			e	
f						-		f	
g							-	g	
h	a ⁻¹	b ⁻¹	c ⁻¹	d ⁻¹	e ⁻¹	f ⁻¹	g ⁻¹	h,h ⁻¹	

Table 5: Inference axioms automatically obtained using the relations from Table 4. A letter indicates an axiom $R_1 \circ R_2 \rightarrow R_3$ by indicating R_3 . An empty cell indicates that R_1 and R_2 do not have compatible domains and ranges; ‘:’ that the composition is prohibited; and ‘-’ that a relation R_3 such that P_{R_3} is consistent with $P_{R_1} \circ P_{R_2}$ could not be found.

yield *Mary is the agent of talking*, and *she has the intention of talking*. These two relations are valid but most probably ignored by a role labeler since *Mary* is not an argument of *talk*.

4 Case Study

In this Section, we apply the CSR algorithm over a set of eight well-known relations. It is out of the scope of this paper to explain in detail the semantics of each relation or their detection. Our goal is to obtain inference axioms and, taking for granted that annotation is available, evaluate their accuracy.

The only requirement for the CSR algorithm is to define semantic relations using the extended definition (Table 4). To define domains and ranges, we use the ontology in Section 4.2. Values for the primitives are assigned manually. The meaning of each relations is as follows:

- CAU(x, y) encodes a relation between two situations, where the existence of y is due to the previous existence of x , e.g., *He [got]_y a bad grade because he [didn’t submit]_x the project*.
- INT(x, y) links an animate concrete object and the situations he wants to become true, e.g., *[Mary]_y would like to [grow]_x bonsais*.
- PRP(x, y) holds between a concept y and its main goal x . Purposes can be defined for situations, e.g., *[pruning]_y allows new [growth]_x*; concrete objects, e.g., *the [garage]_y is used for [storage]_x*; or abstract objects, e.g., *[language]_y is used to [communicate]_x*.
- AGT(x, y) links a situation y and its intentional doer x , e.g., *[Mary]_x [went]_y to Paris*. x is restricted to animate concrete objects.
- MNR(x, y) holds between the mode, way, style or

fashion x in which a situation y happened. x can be a state, e.g., *[walking]_y [holding]_x hands*; abstract objects, e.g., *[die]_y [with pain]_x*; or qualities, e.g. *[fast]_x [delivery]_y*.

- AT-L(x, y) defines the spatial context y of an object or situation x , e.g., *He [went]_x [to Cancun]_y*, *[The car]_x is [in the garage]_y*.
- AT-T(x, y) links an object or situation x , with its temporal information y , e.g., *He [went]_x [yesterday]_y*, *[20th century]_y [sculptures]_x*.
- SYN(x, y) can be defined between any two entities and holds when both arguments are semantically equivalent, e.g., *SYN(dozen, twelve)*.

4.1 Inference Axioms Automatically Obtained

After applying the CSR algorithm over the relations in Table 4, we obtain 78 unique inference axioms (Table 5). Each sub table must be indexed with the first and second premises as row and column respectively. The table on the left summarizes axioms $R_1 \circ R_2 \rightarrow R_3$ and $R_2 \circ R_1 \rightarrow R_3$, the one in the middle axiom $R_1^{-1} \circ R_2 \rightarrow R_3$ and the one on the right axiom $R_2 \circ R_1^{-1} \rightarrow R_3$.

The CSR algorithm identifies several correct axioms and accurately marks as prohibited several combinations that would lead to wrong inferences:

- For CAUSE, the inherent transitivity is detected ($a \circ a \rightarrow a$). Also, no relation is inferred between two different effects of the same cause ($a^{-1} \circ a \rightarrow :$) and between two causes of the same effect ($a \circ a^{-1} \rightarrow :$).
- The location and temporal information of concept y is inherited by its cause, intention, purpose, agent and manner (sub table on the left, f and g columns).

- As expected, axioms involving SYNONYMY as one of their premises yield the other premise as their conclusion (all sub tables).
- The AGENT of y is inherited by its causes, purposes and manners (d row, sub table on the right). In all examples below, $AGT(x, y)$ holds, and we infer $AGT(x, z)$ after composing it with R_2 : (1) $[He]_x [went]_y$ after $[reading]_z$ a good review, R_2 : $CAU^{-1}(y, z)$; (2) $[They]_x [went]_y$ to $[talk]_z$ about it, R_2 : $PRP^{-1}(y, z)$; and (3) $[They]_x [were walking]_y [holding]_z$ hands, R_2 : $MNR^{-1}(y, z)$
An AGENT for a situation y is also inherited by its effects, and the situations that have y as their manner or purpose (d row, sub table on the left).
- A concept intends the effects of its intentions and purposes ($b^{-1} \circ a \rightarrow b^{-1}$, $c^{-1} \circ a \rightarrow b^{-1}$). For example, $[I]_x$ printed the document to $[read]_y$ and $[learn]_z$ the contents; $INT^{-1}(I, read) \circ CAU(read, learn) \rightarrow INT^{-1}(I, learn)$.

It is important to note that domain and range restrictions are not sufficient to identify inference axioms; they only filter out pairs of not compatible relations. The algebra to compose primitives is used to detect prohibited combinations of relations based on semantic grounds and identify the conclusion of composing them. Without primitives, the cells in Table 5 would be either empty (marking the pair as not compatible) or would simply indicate that the pair has compatible domain and range (without identifying the conclusion).

Table 5 summarizes 136 unique pairs of premises (recall $R_i \circ R_j = (R_j^{-1} \circ R_i^{-1})^{-1}$). Domain and range restrictions mark 39 (28.7%) as not compatible. The algebra labels 12 pairs as prohibited (8.8%, [12.4% of the compatible pairs]) and is unable to find a conclusion 14 times (10.3%, [14.4%]). Finally, conclusions are found for 71 pairs (52.2%, [73.2%]). Since more than one conclusion might be detected for the same pair of premises, 78 inference axioms are ultimately identified.

4.2 Ontology

In order to define domains and ranges, we use a simplified version of the ontology presented in (Helbig, 2005). We find enough to contemplate only seven base classes: *ev*, *st*, *co*, *aco*, *ao*, *loc* and *tmp*. Entities (*ent*) refer to any concept and are divided into situations (*si*), objects (*o*) and descriptors (*des*).

- Situations are anything that happens at a time and place and are divided into events (*ev*) and states (*st*). Events imply a change in the status of other entities (e.g., *grow*, *conference*); states do not (e.g., *be standing*, *account for 10%*).
- Objects can be either concrete (*co*, palpable, tangible, e.g., *table*, *keyboard*) or abstract (*ao*, intangible, product of human reasoning, e.g., *disease*, *weight*). Concrete objects can be further classified as animate (*aco*) if they have life, vigor or spirit (e.g. *John*, *cat*).
- Descriptors state properties about the local (*loc*, e.g., *by the table*, *in the box*) or temporal (*tmp*, e.g., *yesterday*, *last month*) context of an entity.

This simplified ontology does not aim at defining domains and ranges for any relation set; it is a simplification to fit the eight relations we work with.

5 Evaluation

An evaluation was performed to estimate the validity of the 78 axioms. Because the number of axioms is large we have focused on a subset of them (Table 6). The 31 axioms having SYN as premise are intuitively correct: since synonymous concepts are interchangeable, given veracious annotation they perform valid inferences.

We use PropBank annotation (Palmer et al., 2005) to instantiate the premises of each axiom. First, all instantiations of axiom $PRP \circ MNR^{-1} \rightarrow MNR^{-1}$ were manually checked. This axiom yields 237 new MANNER, 189 of which are valid (Accuracy 0.80).

Second, we evaluated axioms 1–7 (Table 6). Since PropBank is a large corpus, we restricted this phase to the first 1,000 sentences in which there is an instantiation of any axiom. These sentences contain 1,412 instantiations and are found in the first 31,450 sentences of PropBank.

Table 6 depicts the total number of instantiations for each axiom and its accuracy (columns 3 and 4). Accuracies range from 0.40 to 0.90, showing that the plausibility of an axiom depends on the axiom. The average accuracy for axioms involving CAU is 0.54 and for axioms involving PRP is 0.87.

Axiom $CAU \circ AGT^{-1} \rightarrow AGT^{-1}$ adds 201 relations, which corresponds to 0.89% in relative terms. Its accuracy is low, 0.40. Other axioms are less productive but have a greater relative impact and accu-

No.	Axiom	no heuristic			with heuristic		
		No. Inst.	Acc.	Produc.	No. Inst.	Acc.	Produc.
1	$\text{CAU} \circ \text{AGT}^{-1} \rightarrow \text{AGT}^{-1}$	201	0.40	0.89%	75	0.67	0.33%
2	$\text{CAU} \circ \text{AT-L} \rightarrow \text{AT-L}$	17	0.82	0.84%	15	0.93	0.74%
3	$\text{CAU} \circ \text{AT-T} \rightarrow \text{AT-T}$	72	0.85	1.25%	69	0.87	1.20%
1-3	$\text{CAU} \circ \text{R}_2 \rightarrow \text{R}_3$	290	0.54	0.96%	159	0.78	0.52%
4	$\text{PRP} \circ \text{AGT}^{-1} \rightarrow \text{AGT}^{-1}$	375	0.89	1.66%	347	0.94	1.54%
5	$\text{PRP} \circ \text{AT-L} \rightarrow \text{AT-L}$	49	0.90	2.42%	48	0.92	2.37%
6	$\text{PRP} \circ \text{AT-T} \rightarrow \text{AT-T}$	138	0.84	2.40%	129	0.88	2.25%
7	$\text{PRP} \circ \text{MNR}^{-1} \rightarrow \text{MNR}^{-1}$	71	0.82	3.21%	70	0.83	3.16%
4-7	$\text{PRP} \circ \text{R}_2 \rightarrow \text{R}_3$	633	0.87	1.95%	594	0.91	1.83%
1-7	All	923	0.77	2.84%	753	0.88	2.32%

Table 6: Axioms used for evaluation, number of instances, accuracy and productivity (i.e., percentage of relations added on top the ones already present). Results are reported with and without the heuristic.




Figure 2: Basic (solid arrows) and inferred relations (discontinuous) from *A half-dozen Soviet space officials, in Tokyo in July for an exhibit, stopped by to see their counterparts at the National ...* (wsj_0405, 1).

racy. For example, axiom $\text{PRP} \circ \text{MNR}^{-1} \rightarrow \text{MNR}^{-1}$, only yields 71 new MNR, and yet it is adding 3.21% in relative terms with an accuracy of 0.82.

Overall, applying the seven axioms adds 923 relations on top of the ones already present (2.84% in relative terms) with an accuracy of 0.77. Figure 2 shows examples of inferences using axioms 1-3.

5.1 Error Analysis

Because of the low accuracy of axiom 1, an error analysis was performed. We found that unlike other axioms, this axiom often yield a relation type that is already present in the semantic representation. Specifically, it often yields $R(x, z)$ when $R(x', z)$ is already known. We use the following heuristic in order to improve accuracy: *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 has increased the accuracy of the inferences at the cost of lowering their productivity. The last three columns in Table 6 show results when using the heuristic.

6 Comparison with Previous Work

There have been many proposals to detect semantic relations from text without composition. Researches have targeted particular relations (e.g., CAUSE (Chang and Choi, 2006; Bethard and Martin, 2008)), relations within noun phrases (Nulty, 2007), named entities (Hirano et al., 2007) or clauses (Szpakowicz et al., 1995). Competitions include (Litkowski, 2004; Carreras and Màrquez, 2005; Girju et al., 2007; Hendrickx et al., 2009).

Two recent efforts (Ruppenhofer et al., 2009; Gerber and Chai, 2010) are similar to CSR in their goal (i.e., extract meaning ignored by current semantic parsers), but completely differ in their means. Their merit relies on annotating and extracting semantic connections not originally contemplated (e.g., between concepts from two different sentences) using an already known and fixed relation set. Unlike CSR, they are dependent on the relation inventory, require annotation and do not reason or manipulate relations. In contrast to all the above references and the state of the art, the proposed framework obtains axioms that take as input semantic relations pro-

duced by others and output more relations: it adds an extra layer of semantics previously ignored.

Previous research has exploited the idea of using semantic primitives to define and classify semantic relations under the names of *relation elements*, *deep structure*, *aspects* and *primitives*. The first attempt on describing semantic relations using primitives was made by Chaffin and Herrmann (1987); they differentiate 31 relations using 30 *relation elements* clustered into five groups (intensional force, dimension, agreement, propositional and part-whole inclusion). Winston et al. (1987) introduce 3 *relation elements* (functional, homeomeric and separable) to distinguish six subtypes of PART-WHOLE. Cohen and Losielle (1988) use the notion of *deep structure* in contrast to the *surface relation* and utilizes two *aspects* (hierarchical and temporal). Huhns and Stephens (1989) consider a set of 10 *primitives*.

In theoretical linguistics, Wierzbicka (1996) introduced the notion of semantic primes to perform linguistic analysis. Dowty (2006) studies compositionality and identifies entailments associated with certain predicates and arguments (Dowty, 2001).

There has not been much work on composing relations in the field of computational linguistics. The term *compositional semantics* is used in conjunction with the principle of compositionality, i.e., the meaning of a complex expression is determined from the meanings of its parts, and the way in which those parts are combined. These approaches are usually formal and use a potentially infinite set of predicates to represent semantics. Ge and Mooney (2009) extracts semantic representations using syntactic structures while Copestake et al. (2001) develops algebras for semantic construction within grammars. Logic approaches include (Lakoff, 1970; Sánchez Valencia, 1991; MacCartney and Manning, 2009). Composition of Semantic Relations is complementary to Compositional Semantics.

Previous research has manually extracted plausible inference axioms for WordNet relations (Harabagiu and Moldovan, 1998) and transformed chains of relations into theoretical axioms (Helbig, 2005). The CSR algorithm proposed here automatically obtains inference axioms.

Composing relations has been proposed before within knowledge bases. Cohen and Losielle (1988) combines a set of nine fairly specific relations (e.g.,

FOCUS-OF, PRODUCT-OF, SETTING-OF). The key to determine plausibility is the *transitivity characteristic* of the aspects: two relations shall not combine if they have contradictory values for any aspect. The first algebra to compose semantic primitives was proposed by Huhns and Stephens (1989). Their relations are not linguistically motivated and ten of them map to some sort of PART-WHOLE (e.g. PIECE-OF, SUBREGION-OF). Unlike (Cohen and Losielle, 1988; Huhns and Stephens, 1989), we use typical relations that encode the semantics of natural language, propose a method to automatically obtain the inverse of a relation and empirically test the validity of the axioms obtained.

7 Conclusions

Going beyond current research, in this paper we investigate the composition of semantic relations. The proposed CSR algorithm obtains inference axioms that take as their input semantic relations and output a relation previously ignored. Regardless of the set of relations and annotation scheme, an additional layer of semantics is created on top of the already existing relations.

An extended definition for semantic relations is proposed, including restrictions on their domains and ranges as well as values for semantic primitives. Primitives indicate if a certain property holds between the arguments of a relation. An algebra for composing semantic primitives is defined, allowing to automatically determine the primitives values for the composition of any two relations.

The CSR algorithm makes use of the extended definition and algebra to discover inference axioms in an unsupervised manner. Its usefulness is shown using a set of eight common relations, obtaining 78 axioms. Empirical evaluation shows the axioms add 2.32% of relations in relative terms with an overall accuracy of 0.88, more than what state-of-the-art semantic parsers achieve.

The framework presented is completely independent of any particular set of relations. Even though different sets may call for different ontologies and primitives, we believe the model is generally applicable; the only requirement is to use the extended definition. This is a novel way of retrieving semantic relations in the field of computational linguistics.

References

- Steven Bethard and James H. Martin. 2008. Learning Semantic Links from a Corpus of Parallel Temporal and Causal Relations. In *Proceedings of ACL-08: HLT, Short Papers*, pages 177–180, Columbus, Ohio.
- Eduardo Blanco and Dan Moldovan. 2011. A Model for Composing Semantic Relations. In *Proceedings of the 9th International Conference on Computational Semantics (IWCS 2011)*, Oxford, UK.
- Xavier Carreras and Lluís Màrquez. 2005. Introduction to the CoNLL-2005 shared task: semantic role labeling. In *CONLL '05: Proceedings of the Ninth Conference on Computational Natural Language Learning*, pages 152–164, Morristown, NJ, USA.
- Roger Chaffin and Douglass J. Herrmann, 1987. *Relation Element Theory: A New Account of the Representation and Processing of Semantic Relations*.
- Du S. Chang and Key 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.
- Timothy Chklovski and Patrick Pantel. 2004. VerbOcean: Mining the Web for Fine-Grained Semantic Verb Relations. In *Proceedings of EMNLP 2004*, pages 33–40, Barcelona, Spain.
- Paul R. Cohen and Cynthia L. Losielle. 1988. Beyond ISA: Structures for Plausible Inference in Semantic Networks. In *Proceedings of the Seventh National conference on Artificial Intelligence*, St. Paul, Minnesota.
- Ann Copestake, Alex Lascarides, and Dan Flickinger. 2001. An Algebra for Semantic Construction in Constraint-based Grammars. In *Proceedings of 39th Annual Meeting of the Association for Computational Linguistics*, pages 140–147, Toulouse, France.
- David D. Dowty. 2001. The Semantic Asymmetry of ‘Argument Alternations’ (and Why it Matters). In Geart van der Meer and Alice G. B. ter Meulen, editors, *Making Sense: From Lexeme to Discourse*, volume 44.
- David Dowty. 2006. Compositionality as an Empirical Problem. In Chris Barker and Polly Jacobson, editors, *Papers from the Brown University Conference on Direct Compositionality*. Oxford University Press.
- Ruifang Ge and Raymond Mooney. 2009. Learning a Compositional Semantic Parser using an Existing Syntactic Parser. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 611–619, Suntec, Singapore.
- Matthew Gerber and Joyce Chai. 2010. Beyond NomBank: A Study of Implicit Arguments for Nominal Predicates. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1583–1592, Uppsala, Sweden.
- Roxana Girju, Preslav Nakov, Vivi Nastase, Stan Szpakowicz, Peter Turney, and Deniz Yuret. 2007. SemEval-2007 Task 04: Classification of Semantic Relations between Nominals. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 13–18, Prague, Czech Republic.
- Sanda Harabagiu and Dan Moldovan. 1998. Knowledge Processing on an Extended WordNet. In Christiane Fellbaum, editor, *WordNet: An Electronic Lexical Database and Some of its Applications.*, chapter 17, pages 684–714. The MIT Press.
- Hermann Helbig. 2005. *Knowledge Representation and the Semantics of Natural Language*. Springer, 1st edition.
- Iris Hendrickx, Su N. Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2009. SemEval-2010 Task 8: Multi-Way Classification of Semantic Relations Between Pairs of Nominals. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions (SEW-2009)*, pages 94–99, Boulder, Colorado.
- Toru Hirano, Yoshihiro Matsuo, and Genichiro Kikui. 2007. Detecting Semantic Relations between Named Entities in Text Using Contextual Features. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, Demo and Poster Sessions*, pages 157–160, Prague, Czech Republic.
- Michael N. Huhns and Larry M. Stephens. 1989. Plausible Inferencing Using Extended Composition. In *IJCAI'89: Proceedings of the 11th international joint conference on Artificial intelligence*, pages 1420–1425, San Francisco, CA, USA.
- George Lakoff. 1970. Linguistics and Natural Logic. 22(1):151–271, December.
- Ken Litkowski. 2004. Senseval-3 task: Automatic labeling of semantic roles. In *Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*, pages 9–12, Barcelona, Spain.
- Bill MacCartney and Christopher D. Manning. 2009. An extended model of natural logic. In *Proceedings of the Eight International Conference on Computational Semantics*, pages 140–156, Tilburg, The Netherlands.
- Paul Nulty. 2007. Semantic Classification of Noun Phrases Using Web Counts and Learning Algorithms. In *Proceedings of the ACL 2007 Student Research Workshop*, pages 79–84, Prague, Czech Republic.

- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106.
- Josef Ruppenhofer, Caroline Sporleder, Roser Morante, Collin Baker, and Martha Palmer. 2009. SemEval-2010 Task 10: Linking Events and Their Participants in Discourse. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions (SEW-2009)*, pages 106–111, Boulder, Colorado.
- Victor Sánchez Valencia. 1991. *Studies on Natural Logic and Categorical Grammar*. Ph.D. thesis, University of Amsterdam.
- Barker Szpakowicz, Ken Barker, and Stan Szpakowicz. 1995. Interactive semantic analysis of Clause-Level Relationships. In *Proceedings of the Second Conference of the Pacific Association for Computational Linguistics*, pages 22–30.
- Anna Wierzbicka. 1996. *Semantics: Primes and Universals*. Oxford University Press, USA.
- Morton E. Winston, Roger Chaffin, and Douglas Herrmann. 1987. A Taxonomy of Part-Whole Relations. *Cognitive Science*, 11(4):417–444.