Composition of Semantic Relations: Model and Applications

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

Abstract

This paper presents a framework for combining semantic relations extracted from text to reveal even more semantics that otherwise would be missed. A set of 26 relations is introduced, with their arguments defined on an ontology of sorts. A semantic parser is used to extract these relations from noun phrases and verb argument structures. The method was successfully used in two applications: rapid customization of semantic relations to arbitrary domains and recognizing entailments.

1 Introduction

Semantic representation of text facilitates inferences, reasoning, and greatly improves the performance of Question Answering, Information Extraction, Machine Translation and other NLP applications. Broadly speaking, semantic relations are unidirectional underlying connections between concepts. For example, the noun phrase *the car engine* encodes a PART-WHOLE relation: the engine is a part of the car.

Semantic relations are the building blocks for creating a semantic structure of a sentence. There is a growing interest in text semantics fueled by the new wave of semantic technologies and ontologies that aim at transforming unstructured text into structured knowledge. More and more enterprises and academic organizations have adopted the World Wide Web Consortium (W3C) Resource Description Framework (RDF) specification as a standard representation of text knowledge. This is based on semantic triples, which can be used to represent semantic relations.

The work reported in this paper aims at extracting as many semantic relations from text as possible. Semantic parsers (SP) extract semantic relations from text. Typically they detect relations between adjacent concepts or verb argument structures, leaving considerable semantics unrevealed. For example, given *John is a rich man*, a typical SP extracts *John is a man* and *man is rich*, but not *John is rich*. The third relation can be extracted by combining the two relations detected by the parser. The observation that combining elementary semantic relations yields more relations is the starting point and the motivation for this work.

2 Related Work

In Computational Linguistics, WordNet (Miller, 1995), FrameNet (Baker et al., 1998) and PropBank (Palmer et al., 2005) are probably the most used semantic resources. Like our approach and unlike PropBank, FrameNet annotates semantics between concepts regardless of their position in a parse tree. Unlike us, they use a predefined set of frames to be filled. PropBank adds semantic annotation on top of the Penn TreeBank and it contains only annotations between a verb and its arguments. Moreover, the semantics of a given label depends on the verb. For example, ARG2 is used for INSTRUMENT and VALUE.

Copious work has been done lately on semantic roles (Màrquez et al., 2008). Approaches to detect semantic relations usually focus on particular lexical and syntactic patterns or kind of arguments. There are both unsupervised (Turney, 2006) and supervised approaches. The SemEval-2007 Task 4 (Girju et al., 2007) focused on relations between nominals. Work has been done on detecting relations between noun phrases (Davidov and Rappoport, 2008; Moldovan et al., 2004), named entities (Hirano et al., 2007), and clauses (Szpakowicz et al., 1995). There have been pro-

Coling 2010: Poster Volume, pages 72–80,
Beijing, August 2010

posals to detect a particular relation, e.g., CAUSE (Chang and Choi, 2006), INTENT (Tatu, 2005) and PART-WHOLE (Girju et al., 2006).

Researchers have also worked on combining semantic relations. Harabagiu and Moldovan (1998) combine WordNet relations and Helbig (2005) transforms chains of relations into theoretical axioms. Some use logic as the underlying formalism (Lakoff, 1970; Sánchez Valencia, 1991), more ideas can be found in (Copestake et al., 2001).

3 Approach

In contrast to First Order Logic used in AI to represent text knowledge, we believe text semantics should be represented using a fixed set of relations. This facilitates a more standard representation and extraction automation which in turn allows reasoning. The fewer the relation types, the easier it is to reason and perform inferences. Thus, a compromise has to be made between having enough relation types to adequately represent text knowledge and yet keeping the number small for making the extraction and manipulation feasible.

The main contributions of this paper are: (i) an extended definition of a set of 26 semantic relations resulted after many iterations and pragmatic considerations; (ii) definition of a semantic calculus, a framework to manipulate and compose semantic relations (CSR); (iii) use of CSR to rapidly customize a set of semantic relations; and (iv) use of CSR to detect entailments. The adoption of CSR to other semantic projects does not require any modification of existing tools while being able to detect relations ignored by such tools.

4 Semantic Relations

Formally, a semantic relation is represented as R(*x*, *y*), where R is the relation type and *x* and *y* the first and second argument. R(*x*, *y*) should be read as *x is* R *of y*. The sentence *"John painted his truck"* yields AGENT(*John*, *painted*), THEME(*his truck*, *painted*) and POSSESSION(*truck*, *John*).

Extended definition Given a semantic relation R, DOMAIN(R) and RANGE(R) are defined as the set of sorts of concepts that can be part of the first and second argument. A semantic relation R(*x*, *y*) is defined by its: (i) relation type R; (ii) DOMAIN(R); and (iii) RANGE(R). Stating restrictions for DOMAIN and RANGE has several advantages: it (i) helps distinguishing between relations, e.g., $[\text{tall}]_{ql}$ and $[\text{John}]_{aco}$ can be linked through VALUE, but not POSSESSION; (ii) helps discarding potential relations that do not hold, e.g., abstract objects do not have INTENT; and (iii) helps combining semantic relations (Section 5).

Ontology of Sorts In order to define DOMAIN(R) and RANGE(R), we use a customized ontology of sorts (Figure 1) modified from (Helbig, 2005). The root corresponds to entities, which refers to *all things about which something can be said*.

Objects can be either concrete or abstract. The former occupy space, are touchable and tangible. The latter are intangible; they are somehow a product of human reasoning. Concrete objects are further divided into animate or inanimate. The former have life, vigor or spirit; the later are dull, without life. Abstract objects are divided into temporal or non temporal. The first corresponds to abstractions regarding points or periods of time (e.g. *July*, *last week*); the second to any other abstraction (e.g. *disease*, *justice*). Abstract objects can be sensually perceived, e.g., *pain*, *odor*.

Situations are anything that happens at a time and place. Simply put, if one can think of the time and location of an entity, it is a situation. Events (e.g. *mix*, *grow*) imply a change in the status of other entities, states (e.g. *standing next to the door*) do not. Situations can be expressed by verbs (e.g. *move*, *print*) or nouns (e.g. *party*, *hurricane*).

Descriptors complement entities by stating properties about their spatial or temporal context. They are composed of an optional non-content word signaling the local or temporal context and another entity. Local descriptors are further composed of a concrete object or situation, e.g., *$[\text{above}]_{prep}$ $[\text{the roof}]_{co}$*; temporal descriptors by a temporal abstract object or situation, e.g., *$[\text{during}]_{prep}$ $[\text{the party}]_{ev}$*. The non-content word signaling the local or temporal context is usually present, but not always, e.g., *"The $[\text{birthplace}]_{ev}$ of his mother is $[\text{Ankara}]_{loc}$"*.

Qualities represent characteristics than can be assigned to entities. They can be quantifiable like *tall* and *heavy*, or unquantifiable like *difficult* and *sleepy*. Quantities represent quantitative characteristics of concepts, e.g., *a few pounds*, *22 yards*.

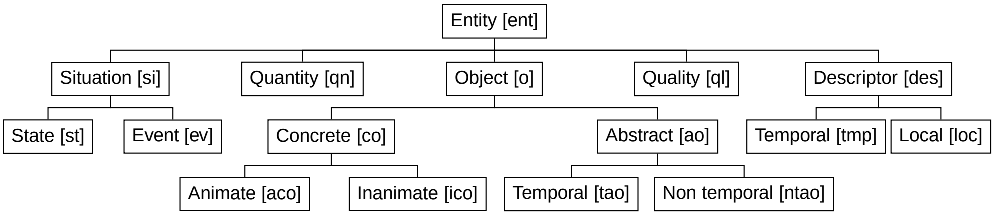

Figure 1: The ontology of sorts of concepts and their acronyms.

Cluster	Relation type	Abr.	Class.	Properties r	s	t	DOMAIN × RANGE	Example
Reason	CAUSE	CAU	iv	-	-	√	[si] × [si]	CAU(*earthquake, tsunami*)
	JUSTIFICATION	JST	iv	-	-	√	[si ∪ ntao] × [si]	JST(*it is forbidden, don't smoke*)
	INFLUENCE	IFL	iv	-	-	√	[si] × [si]	IFL(*missing classes, poor grade*)
Goal	INTENT	INT	i	-	-	-	[si] × [aco]	INT(*teach, professor*)
	PURPOSE	PRP	v	-	-	√	[si ∪ ntao] × [si ∪ co ∪ ntao]	PRP(*storage, garage*)
Object modifiers	VALUE	VAL	v	-	-	-	[ql] × [o ∪ si]	VAL(*smart, kids*)
	SOURCE	SRC	ii	-	-	√	[loc ∪ ql ∪ ntao ∪ ico] × [o]	SRC(*Mexican, students*)
Syntactic subjects	AGENT	AGT	iii	-	-	-	[aco] × [si]	AGT(*John, bought*)
	EXPERIENCER	EXP	iii	-	-	-	[o] × [si]	EXP(*John, heard*)
	INSTRUMENT	INS	iii	-	-	-	[co ∪ ntao] × [si]	INS(*the hammer, broke*)
Direct objects	THEME	THM	iii	-	-	-	[o] × [ev]	THM(*a car, bought*)
	TOPIC	TPC	iii	-	-	-	[o ∪ si] × [ev]	TPC(*flowers, gave*)
	STIMULUS	STI	iii	-	-	-	[o] × [ev]	STI(*the train, heard*)
Association	ASSOCIATION	ASO	v	√	√	√	[ent] × [ent]	ASO(*fork, knife*)
	KINSHIP	KIN	ii	√	√	√	[aco] × [aco]	KIN(*John, his wife*)
None	IS-A	ISA	ii	-	-	√	[o] × [o]	ISA(*gas guzzler, car*)
	PART-WHOLE	PW	ii	-	-	*	[o] × [o] ∪ [l] × [l] ∪ [t] × [t]	PW(*engine, car*)
	MAKE	MAK	ii	-	-	-	[co ∪ ntao] × [co ∪ ntao]	MAK(*cars, BMW*)
	POSSESSION	POS	ii	-	-	√	[co] × [co]	POS(*Ford F-150, John*)
	MANNER	MNR	iii	-	-	-	[ql ∪ st ∪ ntao] × [si]	MNR(*quick, delivery*)
	RECIPIENT	RCP	iii	-	-	-	[co] × [ev]	RCP(*Mary, gave*)
	SYNONYMY	SYN	v	√	√	√	[ent] × [ent]	SYN(*a dozen, twelve*)
	AT-LOCATION	AT-L	v	√	-	*	[o ∪ si] × [loc]	AT-L(*party, John's house*)
	AT-TIME	AT-T	v	√	-	*	[o ∪ si] × [tmp]	AT-T(*party, last Saturday*)
	PROPERTY	PRO	v	-	-	-	[ntao] × [o ∪ si]	PRO(*height, John*)
	QUANTIFICATION	QNT	v	-	-	-	[qn] × [si ∪ o]	QNT(*a dozen, eggs*)

Table 1: The set of 26 relations clustered and classified with their properties (reflexive, symmetric, transitive) and examples. An asterisk indicates that the property holds under certain conditions.

4.1 Semantic Relation Types

This work focuses on the set of 26 semantic relations depicted in Table 1. We found this set specific enough to capture the most frequent semantics of text without bringing unnecessary overspecialization. The set is inspired by several previous proposals. Fillmore introduced the notion of *case frames* and proposed a set of nine roles: AGENT, EXPERIENCER, INSTRUMENT, OBJECT, SOURCE, GOAL, LOCATION, TIME and PATH (Fillmore, 1971). Fillmore's work was extended to FrameNet (Baker et al., 1998). PropBank (Palmer et al., 2005) annotates a set of 17 semantic roles in a per-verb basis.

We aim to encode relations not only between a verb and its arguments, but also between and within noun phrases and adjective phrases. Therefore, more relations are added to the set. It includes relations present in WordNet (Miller, 1995), such as IS-A, PART-WHOLE and CAUSE. Szpakowicz et al. (1995) proposed a set of nine relations and Turney (2006) a set of five. Rosario and Hearst (2004) proposed a set of 38 relations including standard case roles and a set of specific relations for medical domain. Helbig (2005) proposed a set of 89 relations, including ANTONYMY and several TEMPORAL relations, e.g. SUCCESSION, EXTENSION, END.

Our set clusters some of the previous proposals (e.g. we only consider AT-TIME) and discards relations proposed elsewhere when they did not occur frequently enough in our experiments. For example, even though ANTONYMY and ENTAILMENT are semantically grounded, they are very infrequent and we do not deal with them. Our pragmatic goal is to capture as many semantics as possible with as few relations as possible. However, we show (Section 7.1) that our set can be easily customized to a specific domain.

The 26 relations are clustered such that relations belonging to the same cluster are close in meaning. Working with clusters is useful because it allows us to: (i) map to other proposed relations, justifying the chosen set of relations; (ii) work with different levels of specificity; and (iii) reason with the relations in a per cluster basis.

The reason cluster includes relations between a concept having a direct impact on another. CAU(*x*, *y*) holds if *y* would not hold if *x* did not happen. JST(*x*, *y*) encodes a moral cause, motive or socially convened norm. If IFL(*x*, *y*), *x* affects the intensity of *y*, but it does not affect its occurrence.

The goal cluster includes INT and PRP. INT(*x*, *y*) encodes intended consequences, which are volitional. PRP(*x*, *y*) is a broader relation and can be defined for inanimate objects.

The object modifiers cluster encodes descriptions of objects and situations: SRC(*x*, *y*) holds if *x* expresses the origin of *y*. VAL(*x*, *y*) holds for any other attribute, e.g. *heavy*, *handsome*.

The syntactic subjects cluster includes relations linking a syntactic subject and a situation. The differences rely on the characteristics of the subject and the connection per se. AGT(*x*, *y*) encodes an intentional doer, *x* must be volitional. If EXP(*x*, *y*), *x* does not change the situation, it only experiences *y*; it does not participate intentionally in *y* either. If INS(*x*, *y*), *x* is used to perform *y*, *x* is a tool or device that facilitates *y*.

The direct objects cluster includes relations encoding syntactic direct objects. THM(*x*, *y*) holds if *x* is affected or directly involved by *y*. TPC(*x*, *y*) holds if *y* is a communication verb, like *talk* and *argue*. STI(*x*, *y*) holds if *y* is a perception verb and *x* a stimulus that makes *y* happen.

The association cluster includes ASO and KIN. ASO is a broad relation between any pair of entities; KIN encodes a relation between relatives.

The rest of the relations do not fall into any cluster. ISA, PW, SYN, AT-L and AT-T have been widely studied in the literature. MAK(*x*, *y*) holds if *y* makes or produces *x*; POS(*x*, *y*) holds if *y* owns *x*; MNR encodes the way a situation occurs. RCP captures the connection between an event and an object which is the receiver of the event. PRO describes links between a situation or object and its characteristics, e.g., *height*, *age*. Values to the characteristics are given through VAL. QNT(*x*, *y*) holds if *y* is quantitatively determined by *x*.

Relations can also be classified depending on the kind of concepts they describe and their *intra* or *inter* nature into: (i) Intra-Object; (ii) Inter-Objects; (iii) Intra-Situation; (iv) Inter-Situations; and (v) for Object and Situation description.

4.2 Detection of Semantic Relations

Relations are extracted by an in-house SP from a wide variety of syntactic realizations. For example, the compound nominal *steel knife* contains PW(*steel*, *knife*), whereas *carving knife* contains PRP(*carving*, *knife*); the genitive *Mary's toy* contains POS(*toy*, *Mary*), whereas *Mary's brother* contains KIN(*brother*, *Mary*), and *eyes of the baby* contains a PW(*eyes*, *baby*). Relations are also extracted from a verb and its arguments (NP verb, verb NP, verb PP, verb ADVP and verb S), adjective phrases and adjective clauses.

The SP first uses a combination of state-of-the-art text processing tools, namely, part-of-speech tagging, named entity recognition, syntactic parsing and word sense disambiguation. After a candidate syntactic pattern has been found, a series of machine learning classifiers are applied to decide if a relation holds. Four different algorithms are used: decision trees, Naive Bayes, SVM and Semantic Scattering combined in a hybrid approach. Some algorithms use a per-relation approach (i.e., decide whether or not a given relation holds) and others a per-pattern approach (i.e., which relation, if any, holds for a particular pattern). Additionally, human-coded rules are used for a few unambiguous cases. The SP participated in the SemEval 2007 Task 4 (Badulescu and Srikanth, 2007).

5 Composition of Semantic Relations

The goal of semantic calculus (SC) is to provide a formal framework for manipulating semantic relations. CSR is a part of this, its goal is to apply *inference axioms* over already identified relations in text in order to infer more relations.

Semantic Calculus: Operators and Properties The *composition operator* is represented by the

$(\text{R}^{-1})^{-1} = \text{R}$
$\text{R}_i \circ \text{R}_j = ({\text{R}_j}^{-1} \circ {\text{R}_i}^{-1})^{-1}$
R^{-1} inherits all the properties of R
$\perp^{-1} = \perp$
$\forall i$: $\perp \bowtie \text{R}_i$
R is reflexive iff $\forall x$: R(x, x)
R is symmetric iff R(x, y) = R(y, x)
R is transitive iff R(x, y) $\circ$ R(y, z) $\rightarrow$ R(x, z)
$\text{R}_i \rhd \text{R}_j \leftrightarrow {\text{R}_i}^{-1} \lhd {\text{R}_j}^{-1}$
$\text{R}_i \bowtie \text{R}_j \leftrightarrow {\text{R}_i}^{-1} \bowtie {\text{R}_j}^{-1}$
If R_i is symmetric and $\text{R}_i \bowtie \text{R}_j$, ${\text{R}_i}^{-1} \bowtie \text{R}_j$
If R_j is symmetric and $\text{R}_i \bowtie \text{R}_j$, $\text{R}_i \bowtie {\text{R}_j}^{-1}$

Table 2: Semantic calculus properties

symbol $\circ$. It combines two relations and yields a third one. Formally, we denote $\text{R}_1 \circ \text{R}_2 \rightarrow \text{R}_3$.

The *inverse* of R is denoted R^{-1} and can be obtained by simply switching its arguments. Given R(x, y), R^{-1}(y, x) always holds. The easiest way to read R^{-1}(y, x) is *x is* R *of y*.

R_1 *left dominates* R_2, denoted by $\text{R}_1 \rhd \text{R}_2$, iff the composition of R_1 and R_2 yields R_1, i.e., $\text{R}_1 \rhd \text{R}_2$ iff $\text{R}_1 \circ \text{R}_2 \rightarrow \text{R}_1$. R_1 *right dominates* R_2, denoted by $\text{R}_1 \lhd \text{R}_2$, iff the composition of R_2 and R_1 yields R_1, i.e., $\text{R}_1 \lhd \text{R}_2$ iff $\text{R}_2 \circ \text{R}_1 \rightarrow \text{R}_1$. R_1 *completely dominates* R_2, denoted by $\text{R}_1 \bowtie \text{R}_2$, iff $\text{R}_1 \rhd \text{R}_2$ and $\text{R}_1 \lhd \text{R}_2$, i.e., $\text{R}_1 \bowtie \text{R}_2$ iff $\text{R}_1 \circ \text{R}_2 \rightarrow \text{R}_1$ and $\text{R}_2 \circ \text{R}_1 \rightarrow \text{R}_1$.

An OTHER ($\perp$) relation holds between x and y if no relation from the given set holds. Formally, $\perp(x, y)$ iff $\neg\exists \text{R}_i$ such that $\text{R}_i(x, y)$.

Using the notation above, the properties depicted in Table 2 follow.

Necessary conditions for Combining Relations Axioms can be defined only for compatible relations as premises. R_1 and R_2 are *compatible* if it is possible, from a theoretical point of view, to apply the composition operator to them. Formally, $\text{RANGE}(\text{R}_1) \cap \text{DOMAIN}(\text{R}_2) \neq \emptyset$

If R_1 and R_2 are compatible but not equal a *restriction* occurs. Let us denote $\text{RANGE}(\text{R}_1) \cap \text{DOMAIN}(\text{R}_2) = I$. A *backward* restriction takes place if $\text{RANGE}(\text{R}_1) \neq I$ and a *forward* restriction if $\text{DOMAIN}(\text{R}_2) \neq I$. In the former case $\text{RANGE}(\text{R}_1)$ is reduced; in the later $\text{DOMAIN}(\text{R}_2)$ is reduced. A forward and backward restriction can be found with the same pair of relations.

It is important to note that two compatible relations may not be the premises for a valid axiom. For example, KIN and AT-L are compatible but do not yield any valid inference.

Another necessary condition for combining two relations R_1(x, y) and R_2(y, z) is that they have to have a common argument, y.

5.1 Unique axioms

An axiom is defined as a set of relations called premises and a conclusion. Given the premises it unequivocally yields a relation that holds as conclusion. The composition operator is the basic way of combining two relations to form an axiom.

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. Applying property $\text{R}_i \circ \text{R}_j = ({\text{R}_j}^{-1} \circ {\text{R}_i}^{-1})^{-1}$, only 10 combinations are unique: (i) 4 combine R_1, R_2 and their inverses; (ii) 3 combine R_1 and ${\text{R}_1}^{-1}$; and (iii) 3 combine R_2 and ${\text{R}_2}^{-1}$. The most interesting axioms fall into category (i), since the other two can be resolved by the transitivity property of a relation and its inverse.

For n relations there are $2n^2 + n$ potential axioms: $\binom{n}{2} \times 4 + 3n = 2 \times n(n-1) + 3n = 2n^2 + n$. For $n = 26$, there are 1300 potential axioms in (i), 820 of which are compatible.

The number can be further reduced. After manual examination of combinations of ASO and KIN with other relations, we conclude that they do not yield any valid inferences, invalidating 150 potential axioms. This is due to the broad meaning of these relations. QNT can be discarded as well, invalidating 45 more potential axioms.

Some axioms can be easily validated. Because synonymous concepts are interchangeable, SYN is easily combined with any other relation: SYN(x, y) $\circ$ R(y, z) $\rightarrow$ R(x, z) and R(x, y) $\circ$ SYN(y, z) $\rightarrow$ R(x, z). Because hyponyms inherit relations from their hypernyms, ISA(x, y) $\circ$ R(y, z) $\rightarrow$ R(x, z) and R(x, y) $\circ$ ISA^{-1}(y, z) $\rightarrow$ R(x, z) hold. These observations allow us to validate 138 of the 625 potential axioms left, still leaving 487.

As noted before, relations belonging to the same cluster tend to behave similarly. This is especially true for the reason and goal clusters due to their semantic motivation. Working with these two clusters instead of the relations brings the

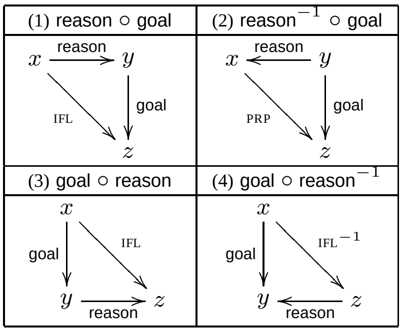

Table 3: The four axioms taking as premises reason and goal clusters. Diagonal arrows indicate inferred relations.

number of axioms to be examined down to 370.

Out of the 370 axioms left, we have extensively analyzed and defined the 35 involving AT-L, the 43 involving reason and the 58 involving goal. Because of space constraints, in this paper we only fully introduce the axioms for reason and goal (Section 6), as well as a variety of axioms useful to recognize textual entailments (Section 7.2).

6 Case Study: Reason and Goal

In this section, we present the four unique axioms for reason and goal relations (Table 3).

(1) **REA(*x*, *y*) ∘ GOA(*y*, *z*) → IFL(*x*, *z*)**: an event is influenced by the reason of its goal.

For example: Bill saves money because he is unemployed; he spends far less than he used to. Therefore, being unemployed can lead to spend far less.

P	REA(*be unemployed, save money*)
	GOA(*save money, spend far less*)
C	IFL(*be unemployed, spend far less*)

(2) **REA^{-1}(*x*, *y*) ∘ GOA(*y*, *z*) → PRP(*x*, *z*)**: events have as their purpose the effects of their goals. This is a strong relation.

For example: Since they have a better view, they can see the mountain range. They cut the tree to have a better view. Therefore, they cut the tree to see the mountain range.

P	REA^{-1}(*see the mountain range, better view*)
	GOA(*better view, cut the tree*)
C	PRP(*see the mountain range, cut the tree*)

Note that possible unintended effects of cutting the tree (e.g. homeowners' association complains) are caused by the event *cut the tree*, not by its effect *get a better view*.

(3) **GOA(*x*, *y*) ∘ REA(*y*, *z*) → IFL(*x*, *z*)**: the goal of an action influences its effects.

For example: John crossed the street carelessly to get there faster. He got run over by a propane truck. Therefore, John got run over by a propane truck influenced by (having the goal of) getting there faster.

P	GOA(*get there faster, crossed carelessly*)
	REA(*crossed carelessly, got run over*)
C	IFL(*get there faster, got run over*)

(4) **GOA(*x*, *y*) ∘ REA^{-1}(*y*, *z*) → IFL^{-1}(*x*, *z*)**. Events influence the goals of its effects.

For example: Jane exercises to lose weight. She exercised because of the good weather. Therefore, good weather helps to lose weight.

P	GOA(*lose weight, exercise*)
	REA^{-1}(*exercise, good weather*)
C	IFL^{-1}(*lose weight, good weather*)

The axioms have been evaluated using manually annotated data. PropBank CAU and PNC are used as reason and goal. Reason annotation is further collected from a corpus which adds causal annotation to the Penn TreeBank (Bethard et al., 2008). A total of 5 and 29 instances for axioms 3 and 4 were found. For all of them, the axioms yield a valid inference. For example, *Buick [approached]$_y$ American express about [a joint promotion]$_x$ because [its card holders generally have a good credit history]$_z$*. PropBank annotation states GOA(*x*, *y*) and REA^{-1}(*y*, *z*), axiom 4 makes the implicit relation IFL^{-1}(*x*, *z*) explicit.

7 Applications and Results

7.1 Customization of Semantic Relations

Problem There is no agreement on a set of relations that best represent text semantics. This is rightfully so since different applications and domains call for different relations. CSR can be used to rapidly customize a set of relations without having to train a new SP or modify any other tool. Given a text, the SP extracts 26 elementary semantic relations. Axioms within the framework of CSR yield n new relations, resulting in a richer semantic representation (Figure 2).

CSR axioms Two ways to get new relations are:

(i) Direct mapping. This is the easiest case and it is equivalent to rename a relation. For example, we can map POS to BELONG or IS-OWNER-OF.

Axiom	Rest. on y	Example
AGT(x, y) $\circ$ THM^{-1}(y, z) $\rightarrow$ ARRESTED(x, z)	*arrested* concept	[Police]$_x$ [apprehended]$_y$ 51 [football fans]$_z$.
THM(x, y) $\circ$ AT-L(y, z) $\rightarrow$ ARRESTED-AT(x, z)	*arrested* concept	Police [apprehended]$_y$ 51 [fans]$_x$ [near the Dome]$_z$.
AGT(x, y) $\circ$ AT-L(y, z) $\rightarrow$ BANKS-AT(x, z)	*banking* activity	[John]$_x$ [withdrew]$_y$ \$20 [at the nearest Chase]$_z$.
POS(x, y) $\circ$ AT-L(y, z) $\rightarrow$ BANKS-AT(x, z)	*account* concept	[John]$_x$ got a [checkbook]$_y$ at [Chase]$_z$.

Table 4: Examples of semantic relation customization using CSR.

Pair	Text T		Hypothesis H	
113	Belknap married and lost his first two wives, Cora LeRoy and Carrie Tomlinson, and married Mrs. John Bower, his second wife's sister.		Belknap was married to Carrie Tomlinson.	
	$T1$	AGT(*Belknap*, *married*)	$H1$	AGT(*Belknap*, *was married*)
	$T2$	THM(*wives*, *married*)	$H2$	THM(*Carrie Tomlinson*, *was married*)
	$T3$	QNT(*first two*, *wives*)		
	$T4$	ISA(*Carrie Tomlinson*, *wives*)		
429	India's yearly pilgrimage to the Ganges river, worshiped by Hindus as the goddess Ganga, is the world's largest gathering of people, …		Ganga is a Hindu goddess.	
	$T1$	AGT(*Hindus*, *worship*)	$H1$	ISA(*Ganga*, *goddess*)
	$T2$	THM(*Ganga*, *worship*)	$H2$	VAL(*Hindu*, *goddess*)
	$T3$	ISA(*Ganga*, *goddess*)		
445	[…] At present day YouTube represents the most popular site sharing on-line video.		YouTube is a video website.	
	$T1$	ISA(*YouTube*, *site*)	$H1$	ISA(*YouTube*, *website*)
	$T2$	EXP(*site*, *sharing*)	$H2$	VAL(*video*, *website*)
	$T3$	THM(*video*, *sharing*)		
716	The Czech and Slovak republics have been unable to agree a political basis for their future coexistence in one country.		The Czech and Slovak republics do not agree to coexist in one country.	
	$T1$	AGT(*The Czech and Slovak republics*, *have been unable to agree*)	$H1$	AGT(*The Czech and Slovak republics*, *do not agree*)
	$T2$	THM(*political basis*, *have been unable to agree*)	$H2$	PRP(*coexist in one country*, *do not agree*)
	$T3$	PRP(*their future coexistence in one country*, *political basis*)		
771	In 2003, Yunus brought the microcredit revolution to the streets of Bangladesh to support more than 50,000 beggars, whom the Grameen Bank respectfully calls Struggling Members.		Yunus supported more than 50,000 Struggling Members.	
	$T1$	AGT(*Yunus*, *brought*)	$H1$	AGT(*Yunus*, *supported*)
	$T2$	PRP(*support*, *brought*)		
	$T3$	RCP(*beggars*, *support*)	$H2$	RCP(*Struggling Members*, *support*)
	$T4$	QNT(*more than 50,000*, *beggars*)	$H3$	QNT(*more than 50,000*, *Struggling Members*)
	$T5$	SYN(*beggars*, *Struggling Members*)		

Table 5: RTE3 examples and their elementary semantic relations (i.e., the ones the SP detects). Only relevant semantic relations for entailment detection are shown for T.

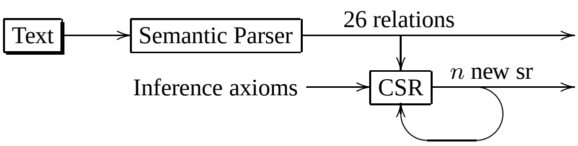

Figure 2: Flowchart for obtaining customized semantic relations using CSR.

(ii) Combinations of two elementary relations yield new specialized relations. In this case, restrictions on the arguments must be fulfilled.

Consider we need the new relation ARRESTED(x, y), which encodes the relation between two animate concrete objects x and y, where x arrested y. We can infer this relation by using the following axiom: AGENT(x, y) $\circ$ THEME^{-1}(y, z) $\rightarrow$ ARRESTED(x, z) provided that y is an *arrested* concept. A simple way of checking if a given concept is of a certain kind is to check WordNet. Collecting all the words belonging the the synset arrest.v.1, we get the following list of *arrested* concepts: *collar*, *nail*, *apprehend*, *pick up*, *nab* and *cop*. Using lexical chains the list could be further improved.

More examples of axioms for generating customized semantic relations are shown in Table 4.

Results Virtually any domain could be covered by applying customization over the set of 26 relations. The set has been successfully customized to a law enforcement domain. Axioms for a total of 37 new relations were defined and implemented. Among others, axioms to infer IS-EMPLOYER, IS-COWORKER, IS-PARAMOUR, IS-INTERPRETER, WAS-ASSASSIN, ATTENDS-SCHOOL-AT, JAILED-AT, COHABITS-WITH, AFFILIATED-TO, MARRIED-TO, RENTED-BY, KIDNAPPED-BY and the relations in Table 4 were defined. Note that a relation can be inferred by several axioms. This customization effort to add 37 new specialized relations took a person only a few days and without modifying the SP.

7.2 Textual Entailment

Problem An application of CSR is recognizing entailments. Given text T and hypothesis H, the task consists on determining whether or not H can be inferred by T (Giampiccolo et al., 2007).

CSR axioms Several examples of the RTE3 challenge can be solved by applying CSR (Table 5). The rest of this section depicts the axioms involved in detecting entailment for each pair.

Pair 113 is a simple one. A perfect match for H in T can be obtained by an axiom reading *all concepts inherit the semantic relations of their hypernyms*. Formally, $\text{ISA}(x, y) \circ \text{THM}(y, z) \rightarrow \text{THM}(x, z)$, $T2$ and $T4$ are the premises and the conclusion matches $H2$. $T1$ matches $H1$.

Pair 429 can be solved by an axiom reading *agents are values for their themes*. Formally, $\text{AGT}(x, y) \circ \text{THM}^{-1}(y, z) \rightarrow \text{VAL}(x, z)$; $T1$ and $T2$ yield VAL(*Hindu*, *Ganga*), which combined with $T3$ results in a match between T and H.

Pair 445 follows a similar pattern, but the way an EXP combines with its THM differs from the way an AGT does. The *theme is a value of the experiencer*, $\text{THM}(x, y) \circ \text{EXP}^{-1}(y, z) \rightarrow \text{VAL}(x, z)$. Given $T2$ and $T3$, the axiom yields $T4$: VAL(*video*, *site*). Assuming that SYN(*site*, *website*), $T1$ and $T4$ match H.

Pair 716 also requires only one inference step. Using $T3$ and $T2$, an axiom reading *situations have as their purpose the purpose of its theme* infers $H2$, yielding a perfect match between T and H. Formally, $\text{PRP}(x, y) \circ \text{THM}(y, z) \rightarrow \text{PRP}(x, z)$.

Pair 771 Using as premises $T1$ and $T2$, an axiom reading *an agent performs the purposes of its actions* infers $H1$. Using $T3$ and $T5$, and $T4$ and $T5$ as premises, an axiom reading *synonymous concepts are interchangeable* infers $H2$ and $H3$, resulting in a perfect match between T and H. Formally, $\text{AGT}(x, y) \circ \text{PRP}^{-1}(y, z) \rightarrow \text{AGT}(x, z)$, $\text{RCP}^{-1}(x, y) \circ \text{SYN}(y, z) \rightarrow \text{RCP}^{-1}(x, z)$ and $\text{QNT}(x, y) \circ \text{SYN}(y, z) \rightarrow \text{QNT}(x, z)$.

Results We conducted two experiments to quantify the impact of CSR in detecting entailments.

First, 60 pairs were randomly selected from the RTE3 challenge and parsed with the SP. 14 of them (23%) could be solved by simply matching the elementary relations in T and H. After applying CSR, 21 more pairs (35%) were solved. Thus, adding CSR on top of the SP clearly improves entailment detection. Out of the 25 pairs not solved, 5 (8%) need coreference resolution and 20 (34%) require commonsense knowledge or fairly complicated reasoning methods (e.g. *a shipwreck is a ship that sank*).

CSR has also been added to a state of the art system for detecting textual entailment (Tatu and Moldovan, 2007). Prior to the addition, the system made 222 errors consisting of 46 false negatives (examples in Table 5) and 176 false positives. CSR was able to correctly solve 18 (39%) of the 46 false negatives.

8 Conclusions

Although the idea of chaining semantic relations has been proposed before, this paper provides a formal framework establishing necessary conditions for composition of semantic relations. The CSR presented here can be used to rapidly customize a set of relations to any arbitrary domain. In addition to the customization of an information extraction tool and recognizing textual entailments, CSR has the potential to contribute to other applications. For example, it can help improve a semantic parser, it can be used to acquire commonsense knowledge axioms and more.

When an axiom that results from combining two relations does not always hold, it may be possible to add constraints that limit the arguments of the premises to only some concepts.

This work stems from the need to automate the extraction of deep semantics from text and representing text as semantic triples. The paper demonstrates that CSR is able to extract more relations than a normal semantic parser would.

References

Badulescu, Adriana and Munirathnam Srikanth. 2007. LCC-SRN: LCC's SRN System for SemEval 2007 Task 4. In *Proceedings of the Fourth International Workshop on Semantic Evaluations*, pages 215–218.

Baker, Collin F., Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the 17th international conference on Computational Linguistics*, Montreal, Canada.

Bethard, Steven, William Corvey, Sara Klingenstein, and James H. Martin. 2008. Building a Corpus of Temporal-Causal Structure. In *Proceedings of the Sixth International Language Resources and Evaluation Conference*, Marrakech, Morocco.

Chang, Du S. 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.

Copestake, Ann, Alex Lascarides, and Dan Flickinger. 2001. An Algebra for Semantic Construction in Constraint-based Grammars. In *Proceedings of 39th Annual Meeting of the ACL*, pages 140–147.

Davidov, Dmitry and Ari Rappoport. 2008. Classification of Semantic Relationships between Nominals Using Pattern Clusters. In *Proceedings of ACL-08: HLT*, pages 227–235, Columbus, Ohio.

Fillmore, Charles J. 1971. Some Problems for Case Grammar. *Monograph Series on Languages and Linguistics*, 24:35–36.

Giampiccolo, Danilo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. The Third PASCAL Recognizing Textual Entailment Challenge. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9.

Girju, Roxana, Adriana Badulescu, and Dan Moldovan. 2006. Automatic Discovery of Part-Whole Relations. *Computational Linguistics*, 32(1):83–135.

Girju, Roxana, 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*, pages 13–18, Prague, Czech Republic.

Harabagiu, Sanda and Dan Moldovan. 1998. Knowledge Processing on an Extended WordNet. In Fellbaum, Christiane, editor, *WordNet: An Electronic Lexical Database and Some of its Applications*, chapter 17, pages 684–714. The MIT Press.

Helbig, Hermann. 2005. *Knowledge Representation and the Semantics of Natural Language*. Springer.

Hirano, Toru, 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 ACL, Demo and Poster Sessions*, pages 157–160.

Lakoff, George. 1970. Linguistics and Natural Logic. *Synthese*, 22(1):151–271.

Màrquez, Lluís, Xavier Carreras, Kenneth C. Litkowski, and Suzanne Stevenson. 2008. Semantic Role Labeling: An Introduction to the Special Issue. *Computational Linguistics*, 34(2):145–159.

Miller, George A. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*, 38:39–41.

Moldovan, Dan, Adriana Badulescu, Marta Tatu, Daniel Antohe, and Roxana Girju. 2004. Models for the Semantic Classification of Noun Phrases. In *HLT-NAACL 2004: Workshop on Computational Lexical Semantics*, pages 60–67.

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

Rosario, Barbara and Marti Hearst. 2004. Classifying Semantic Relations in Bioscience Texts. In *Proc. of the 42nd Meeting of the ACL*, pages 430–437.

Sánchez Valencia, Victor. 1991. *Studies on Natural Logic and Categorial Grammar*. Ph.D. thesis, University of Amsterdam.

Szpakowicz, Barker, Ken Barker, and Stan Szpakowicz. 1995. Interactive semantic analysis of Clause-Level Relationships. In *Proceedings of the Second Conference of the Pacific ACL*, pages 22–30.

Tatu, Marta and Dan Moldovan. 2007. COGEX at RTE 3. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 22–27, Prague, Czech Republic.

Tatu, Marta. 2005. Automatic Discovery of Intentions in Text and its Application to Question Answering. In *Proceedings of the ACL Student Research Workshop*, pages 31–36, Ann Arbor, Michigan.

Turney, Peter D. 2006. Expressing Implicit Semantic Relations without Supervision. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the ACL*, pages 313–320, Sydney, Australia.