Natural Language Engineering **18** (3): 343–374. © Cambridge University Press 2011
doi:10.1017/S1351324911000222

Semantic composition of AT-LOCATION relation with other relations

HAKKI C. CANKAYA[1], EDUARDO BLANCO[2] and DAN MOLDOVAN[2]

[1]*Department of Computer Engineering, Izmir University of Economics, Izmir 35330, Turkey*
email: hakki.cankaya@ieu.edu.tr
[2]*Human Language Technology Research Institute, University of Texas at Dallas, Richardson, TX 75080, USA*
emails: {eduardo, moldovan}@hlt.utdallas.edu

(*Received 13 August 2010; revised 22 April 2011; accepted 6 June 2011; first published online 18 August 2011*)

Abstract

This paper presents a method for the composition of AT-LOCATION with other semantic relations. The method is based on inference axioms that combine two semantic relations yielding another relation that otherwise is not expressed. An experimental study conducted on PropBank, WordNet, and eXtended WordNet shows that inferences have high accuracy. The method is applicable to combining other semantic relations and it is beneficial to many semantically intense applications.

1 Introduction

Semantic representation of text is of great importance for natural language processing (NLP), since many applications depend on it. For example, semantic representation facilitates inferences and reasoning, which in turn may impact Question Answering, Information Extraction, Text Summarization, and other applications. In this work, we distinguish between semantic roles and semantic relations. Semantic roles depict associations between verbs and their arguments (Baker, Fillmore and Lowe 1998; Pradhan *et al.* 2004; Màrquez *et al.* 2008). Semantic relations depict semantic associations between concepts found in compound nouns, noun phrases, verb phrases, clauses, in addition to verb argument structures. For example, the compound noun *car door* encodes a part-whole relation: the *door* is part of the *car*. Similarly, the statement *His extracurricular activities lowered his grades* carries a *causal* relation between the *activities* and the *lowered his grades*. Detecting this relation would help answer a number of questions, including *Why did he have low grades?*

Semantic relations are used to build the semantic structure of a sentence, which in turn helps transform unstructured text into structured knowledge. The interest level is increasing in such semantic technologies and enabling ontologies.

A typical approach for extracting semantic relations between concepts is based on some syntactic methods (rules) and/or patterns that are likely to encode some specific relations (Moldovan *et al.* 2004; Chang and Choi 2006; Girju *et al.* 2007; Bethard and Martin 2008; Srikumar *et al.* 2008). In general, this approach is mostly valid, but overlooks some relations between concepts that do not follow a particular pattern or fails to disambiguate relations that follow the same or similar patterns (Mihalcea and Shi 2005). Moreover, the approach is not able to retrieve the potential relations between concepts that are not within some close proximity from each other. For example, from the sentence '*John's office is in Carrington Park, located in Highland Lakes district*' a semantic parser would likely extract that *John's office* is located at *Carrington Park* and *Carrington Park* is located at *Highland Lakes*. However, it is unlikely that it will extract the simple information that *John's office is in Highland Lakes*.

This observation, where more semantics can be extracted from text in addition to what a semantic parser extracts, was the starting point of our investigation. We propose a method for augmenting already extracted semantics by combining multiple relations over common concepts. We call this composition of semantic relations (CSR). Pairs of relations can be formed and their semantic combinations studied over broad range of concepts and contexts.

In this paper, we focus on the composition of AT-LOCATION relation with other semantic relations in a sample semantic relations set. The results show the proposed method can increase the total number of instances of semantic relations extracted by a semantic parser. We anticipate that such significant improvements in revealing additional semantics will benefit many semantically based applications.

The CSR proposed here uses axioms that conclude a relation from the original set. The axioms have properties and necessary conditions to apply. There are four main benefits of CSR: (i) the method detects relations that escape a regular Semantic Parser (SP); (ii) the method is able to detect relations between concepts that are far apart from each other; (iii) the method does not require modifications and/or configuration change to existing tools; and (iv) the method has a high precision/accuracy.

Furthermore, CSR has been proven useful in several pragmatic applications, such as the customization of a set of relations to a specific domain, recognizing textual entailment (TE), and commonsense knowledge acquisition.

2 Related previous works

The semantics of relations between concepts has been studied in philosophy, psychology, and anthropology as well as in linguistics. Philosophers have considered the intrinsic physical meaning of space and location as well as how they relate to other concepts, their importance in human and animal life, and impact on several other fields of science (Halsbury 1962; Whiteman 1967).

Space and location are very tangible and intuitive notions. Children acquire this knowledge in their early stages of development, partly due to the important survival instinct (e.g., *finding the way back home for safety*) (Piaget 1952; Heft and Wohwill 1987; Golledge *et al.* 1992).

Even creatures without the ability to speak are able to use space and location to reason and infer information. It has been claimed that humans also use spatial reasoning more effectively and easily than logical reasoning (Jackendoff 1983; Jackendoff 1990). Even though spatial cognition is important for the survival of human beings, the amount of spatial reckoning and description may vary across cultures (Levinson 1996). Some cultures may provide fewer locations in their descriptions, assuming that they are obvious. The same behavior can be observed in some other cultures due to ethnocentrisms (Levinson 1996). For example, directions to the same location could be given differently in different cultures because of the difference in linguistic resources. In these cases, automatically augmenting location knowledge from what is already revealed (spatial knowledge) becomes even more important. The AT-LOCATION semantic relation is used to express the spatial characterization of concepts. Augmenting AT-LOCATION, then, would carry the same importance.

There have been numerous studies on semantic relations (Szpakowicz and Barker 1995; Moldovan *et al.* 2004; Turney 2006; Srikumar *et al.* 2008; Beamer and Girju 2009) that propose different sets of relations. The question remains still unanswered about the cardinality of the set. One could theoretically choose to use a set of only two relations, e.g., RELATED and NOT-RELATED. On the other hand, there are proposals with hundreds of very specific relations, such as FROM-AIRPORT and TO-AIRPORT in an investigative context. Helbig (Helbig 2005), in his detailed study, proposed a set of eighty-nine relations without addressing how to extract them from text, which is a different problem. His relations are fine-grained and his set includes multiple relations for expressing common concepts such as time (e.g., SUCCESSION, DURATION, END). Most researchers use no more than a couple of dozens of semantic relations for a generic use. One can argue that there is no exact number nor types of semantic relations. It is simply the application that ultimately dictates a suitable set of relations. In this study, we use a set of twenty-six relations.

In computational linguistics, the most used resources for semantics are (i) WordNet (Miller 1995); (ii) FrameNet (Baker *et al.* 1998); and (iii) PropBank (Kingsbury, Palmer and Marcus 2002; Palmer, Gildea and Kingsbury 2005).

WordNet is a lexical database in which nouns, verbs, adjectives, and adverbs are grouped into synsets. Each synset expresses a distinct concept and is linked to other synsets by semantic and lexical relations. For example, the relations for nouns include ANTONYM, HYPERNYM, and three types of PART-WHOLE (PART-OF, SUBSTANCE-OF, and MEMBER-OF). It has also relations for verbs including ANTONYM, HYPERNYM, ENTAILMENT, and CAUSE.

FrameNet[1] is based on Fillmore's work and defines *semantic frames* and *frame elements* as their participants. Then, the lexical units that invoke the frame are selected. For example, the frame *body movement* may be invoked by a set of verbs

[1] http://framenet.icsi.berkeley.edu/

including *arch*, *bat*, and *wave*, etc. In this example the core frame elements are the *agent* and the *body part* that is to be moved. An instantiation of the frame is $[Mary]_{agent}$ waved $[her\ hand]_{bodypart}$. FrameNet also annotates semantics between any two concepts in the sentence and uses a set of predefined frames to be filled or instantiated.

PropBank uses the Penn TreeBank (Marcus, Santorini and Marcinkiewicz 1994) and adds semantic annotation on top by using seventeen semantic roles. It labels the relations between a verb and its arguments. Automatic detection of semantic roles have drown considerable attention (Blaheta and Charniak 2000; Chen and Rambow 2003; Carreras and Màrquez 2004; Pradhan *et al.* 2004; Giuglea and Moschitti 2006; Màrquez *et al.* 2008), and it has been proven to be useful for several applications. Among others, these are coreference resolution (Ponzetto and Strube 2006), word sense disambiguation (Dang and Palmer 2005), and machine translation (Wu and Fung 2009).

There have been some recent efforts on semantic relation extraction. Probabilistic models with data mining techniques are used for both implicit relation extraction and relational pattern recognition (Culotta, McCallum and Betz 2006). The SemEval-2007 Task 04 (Girju *et al.* 2007) and SemEval-2010 Task 08 (Hendrickx *et al.* 2010) aimed at relations between nominals. The former examined the problem of deciding whether, for a set of seven, any given relation holds (binary classification). The latter considered the more complex problem of which relations hold for a set of nine. The main difference is that, rather than text, our approach takes relations previously extracted, and outputs more relations. It does not require modification in any tool.

Researchers have also studied how to connect concepts following chain of semantic relations. Miller and Johnson-Laird have presented the idea of composing relations that include location in their book titled 'Language and Perception' (Miller and Johnson-Laird 1976). Harabagiu (Harabagiu 1998) uses *lexico-semantic* connections between concepts from WordNet to derive metonymic coercions. Harabagiu and Moldovan (Harabagiu and Moldovan 1998) combine WordNet relations by using part-of-speech (POS) tags. From theoretical point of view, Kenny and Helbig (Kenny 1966; Helbig 2005) offer some inference axioms by using chains of relations. Winston *et al.* (Winston, Chaffin and Herrmann 1987) proposed a larger taxonomy for meronymy and studied the transitivity of meronymy among inclusion relations. The term composition of semantic relations is used in conjunction with the principle of compositionality, that is, the meaning of a complex expression is determined from the meanings of its parts, and the way in which those parts are combined.

There are few studies which focus on a particular semantic relation and its combination of other semantic information. Tatu (Tatu 2007) studies INTENTION in this respect. To the best of our knowledge, the different ways of combining relations have not been deeply studied. In our study, we focus on a formal method for the composition of semantic relations and use the method for augmenting spatial knowledge by doing complete composition analysis of AT-LOCATION relation within a given set of semantic relations.

3 Semantic relations

3.1 Definition

Semantic relations represent the implicit associations between two concepts expressed by words or phrases. Formally, a semantic relation is denoted by R(x, y), where R is the **relation type** and x and y are the **first and second arguments**, respectively. R(x, y) should be interpreted as *x is* R *of y*. The relation type dictates the nature of the association between the two arguments.

For each relation type, the first and second arguments belong to some restricted sets of concepts. In order to accommodate this restriction, we define these sets as DOMAIN(R) for the first argument and RANGE(R) for the second, namely, x ∈ DOMAIN(R) and y ∈ RANGE(R). For example, suppose AGENT(x, y) encodes x as the agent of activity y. Then, we restrict the first argument x of the relation AGENT only for objects that can be an AGENT, filtering out other objects.

A semantic relation R(x, y) is formally defined by stating the following: (i) relation type R: R(x, y) holds if *x is* R *of y*; (ii) DOMAIN(R); and (iii) RANGE(R). This **extended formal definition** incorporates the sets of concepts that can appear as the first and second argument for R by DOMAIN(R) and RANGE(R).

In order for a relation R(x, y) to hold, the first requirement is its domain and range compatibility. Formally, R(x, y) is considered domain and range compatible iff: $x \in$ DOMAIN(R) and $y \in$ RANGE(R). The second requirement is that the relation type encodes the association for the arguments. If the domain and range compatibility fails, then there is no point checking the relation type for the correct association.

The classification of concepts for DOMAIN and RANGE of relations has been attempted by using POS tags (Harabagiu and Moldovan 1998). Although this has brought some restrictions for DOMAIN and RANGE, it was not sufficient and was too coarse of a classification. For example, say an event can be expressed by a verb (*read*, *play*) or a noun (*report*, *card*). Consider that AGENT holds between a noun and a verb, noun restricting the agent and verb restricting the activity. With these restrictions, one could state that AGENT(*pen*, *read*) holds, which clearly should not because the domain of AGENT relation is the set of animate objects.

Recently, a semantic classification of entities has been proposed in order to formally define the sorts of concepts for DOMAIN and RANGE of semantic relations (Helbig 2005). This classification is performed by an **ontology of entities**, which defines a hierarchy of **sorts of concepts** following a semantic criteria. For example, AGENT holds between animate objects (*Mary*, *John*, etc.) and situations (*read*, *run*, etc.). Therefore, it does not make sense to use inanimate objects (*table*, *book*) or abstract objects (*Sunday*, *thought*) for the DOMAIN of AGENT.

Defining the sorts of concepts for DOMAIN and RANGE can have multiple advantages. For one, it can help differentiate between relations that are semantically close but have a little different concepts in their DOMAIN and RANGE. For example, INTENT and PURPOSE are very close in meaning and they both belong to the same semantic relations cluster (explained later in more detail). However, their DOMAIN and RANGE have different sorts of concepts. Only animate objects, aco in Figure 1, can have an INTENT (RANGE(INTENT)=[aco]), while this restriction does not exist for PURPOSE

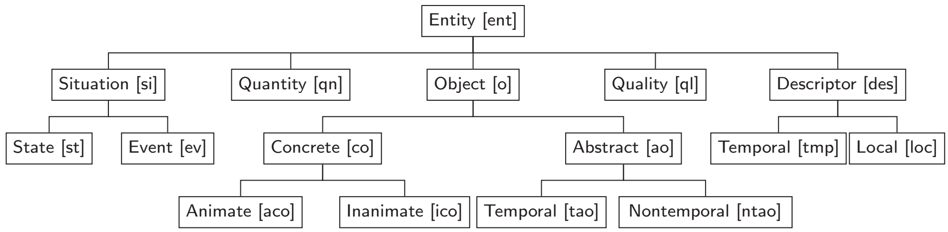

Fig. 1. The ontology of sorts of concepts and their acronyms.

(Range(purpose)=[si ∪ co ∪ ntao]). Therefore, the distinction between intent and purpose can be made easier by taking into account the *sorts of concepts* involved.

Another advantage is to facilitate post filtering to discard the semantic relations that do not comply with the domain and range constraints. A generic semantic parser today determines the semantic relations in a text. Most parsers (Badulescu and Srikanth 2007) achieve a relatively modest performance, typically yielding high recall but low precision. For example, consider the following sentences: (i) *"John works at Starbucks for fun"*; and (ii) *"John got the credit for himself"*. Since both examples have the same syntactic pattern, a semantic parser would probably extract the following two relations: purpose(*fun*, *working at Starbucks*) and purpose(*himself*, *got credit*). Utilizing the defined sorts of concepts for the Domain of purpose relation (Domain(purpose) = [si ∪ ntao]), it would be easy to filter out the incorrect semantic relation, as in this case *himself*, which is an animate concrete object, does not belong to the Domain of purpose relation (aco ∉ Domain(purpose)) and should not be a purpose of *"getting credit"*.

Sorts of concepts would also help define the rules for combining semantic relations. To combine two semantic relations, they should share at least some sorts of concepts in their Domain and Range. Therefore, by simply checking domain and range compatibilities, the possible valid combinations of semantic relations can be analyzed to obtain *inference axioms*.

3.2 The ontology

We use an ontology of entities in the definitions of semantic relations. For the upper ontology of entities, we have adapted the basic structure of an earlier proposal (Helbig 2005), with some modifications. The ontology of this earlier proposal covers the set of its semantic relations (Helbig 2005), which is a larger set than ours. Thus, we simplified the ontology for our set of relations. We also introduced extra sorts of concepts to make the differentiation between some of our relations that are close in meaning.

Figure 1 plots the modified ontology. The root corresponds to entity, which refers to *all things about which something can be said* (Helbig 2005). At the highest level of the hierarchy there are objects, situations, descriptors, quantities, and qualities.

Objects can be either concrete or abstract. The former occupy space, and are touchable and tangible. The latter are intangible, they are somehow product of human reasoning.

Concrete objects are further divided in animate and inanimate. The former have life, vigor, or spirit; the latter are dull, not enlivened. For example, *John* and *Mary's greyhound* are animate; *paper* and *keyboard* are inanimate.

Abstract objects are divided into temporal and nontemporal. The first correspond to abstractions regarding points or periods of time (*July*, *yesterday*, *last month*); the second, to any other abstraction (*disease*, *weight*, *justice*). Note that abstract objects can also be sensually perceived. For example, *pain*, *odor*, *fear*, and *love* are abstract and can be sensed.

Situation is anything that happens at a particular time and place. If a situation implies a change in the status of another entity, then it is called an event (*mix*, *grow*), otherwise a state (*standing*). Situations can be expressed by verbs (*move*, *print*) or nouns (*party*, *hurricane*, *conference*).

Descriptors complement entities by stating the properties of their spatial or temporal context. Therefore, they are divided into local/spatial (*Texas*, *Japan*, *etc.*) and temporal descriptors (*Week*, *July*, *etc.*).

Qualities represent characteristics that can be assigned to entities. These may be quantifiable, such as *tall* and *heavy*, or unquantifiable such as *difficult* and *sleepy*. For example, MANNER semantic relation uses Qualities to describe the way a situation occurs, MANNER(*difficult*, *talk*).

Quantities represent quantitative characteristics of concepts, for example, *a few pounds*, *twenty-two yards*, or *two dozens*.

3.3 Properties

1.Inverse: The inverse of R is denoted R^{-1} and can be obtained by simply switching the arguments of R. Given $R(x, y)$, $R^{-1}(y, x)$ always holds. Formally, $R(x, y) \leftrightarrow R^{-1}(y, x)$ and $(R(x, y))^{-1} = R^{-1}(y, x)$. The easiest way to read $R^{-1}(y, x)$ is *x is* R *of y* and $(R^{-1})^{-1} = R$.

2.Symmetry: R is symmetric iff $R(x, y) = R(y, x)$.

3.4 A list of semantic relations

In this section, we depict a set of twenty-six relations we work with. We believe that the set is specific enough to cover the most frequent semantics of text without overspecialization. Several previous proposals had inspired the set. Fillmore (Fillmore 1971) introduced a set of nine roles with the notion of case frames. Those include AGENT, EXPERIENCER, INTENT, OBJECT, SOURCE, GOAL, AT-LOCATION, AT-TIME, PATH. PropBank annotates seventeen semantic roles in connection with verbs (Palmer *et al.* 2005). In our set, we considered relations not only between a verb and its arguments, but also between and within noun phrases and adjective phrases. We have used some relations from WordNet (Miller 1995), such as IS-A, PART-WHOLE, and CAUSE, but excluded some of the relations which did not occur frequently enough in our experiments (e.g., ANTONYMY, ENTAILMENT). While building the set, we have maintained the goal of capturing as much semantics as possible with as few relations as possible. The set and its similar versions have been used in Moldovan *et al.* (2004) and Tatu and Moldovan (2007). Relations are **clustered** in a semantic way, that is, relations belonging to the same cluster are close in meaning. Working with clusters is useful for several reasons:

Helping to justify the chosen set of relations: Relations are clustered into general meanings. Most of the previous works share with us a general meaning for the relations, but the details about each particular relation may be different. Having a common ground helps comparing and mapping our relations to other proposals and previous annotation efforts.

Allowing to work with different levels of specificity: In some cases, it may be enough to just know the general meaning of the connection between the concepts. For example, if we are asked whether a certain situation s_1 has a direct impact on another situation s_2, it is good enough to check if *any* of the relations belonging to the cluster reason holds between the two situations.

Allowing to reason with the relations per cluster basis: Since relations belonging to the same cluster are close in meaning, they tend to behave similarly. This fact allows the creation of inference axioms that involve clusters instead of single relations. By doing so, the number of potential axioms is decreased.

Table 1 tabulates the set of twenty-six relations with their abbreviations and clusters.

The reason cluster includes relations between a concept having a **direct impact** on another:

- CAUSE(x, y) holds if y would not hold if x did not happen. This relation can only hold between situations, since objects can be *generated or created* (encoded by MAK), but *not caused*. For example, *He [got a bad grade]$^{y}_{ev}$ because he didn't [submit the project]$^{x}_{ev}$*; CAU(*didn't submit the project, got a bad grade*).
- JUSTIFICATION(x, y) is very close to CAU, it encodes a moral cause or motive. If JST(x, y), y would not hold if x did not happen and x is a moral reason or socially convened norm. The distinction between CAU and JST depends on the nature of x. For example, *They [do not smoke in the hall]$^{y}_{st}$ because it [is forbidden]$^{x}_{st}$*; JST(*is forbidden, do not smoke in the hall*).
- INFLUENCE(x, y) encodes a weaker relation than CAU(x, y). If IFL(x, y), x affects the intensity of y, but it does not affect the occurrence. An event may have several influencers. For example, *[Exercising regularly]$^{x}_{st}$ can have an affect on [living a healthier life]$^{y}_{st}$*; IFL(*Exercising regularly, living a healthier life*).

The goal cluster includes the relations INTENT and PURPOSE, which are very close and sometimes it is difficult to distinguish between them. One *intends* to do something for a *purpose*, for example, *"Mary intends to buy a dress to look pretty"*. Both relations bring **uncertainty**, since having an intention or purpose does not guarantee that it will hold:

- INTENT(x, y) encodes intended consequences, which are volitional. Therefore, the range is restricted to animate concrete objects and it does not make sense to consider intentions for abstract objects or other sorts. Situations do not

Table 1. *Semantic clusters and twenty-six relations*

Cluster	Relation type	Abbreviation	DOMAIN × RANGE
Reason	CAUSE	CAU	[si] × [si]
	JUSTIFICATION	JST	[si ∪ ntao] × [si]
	INFLUENCE	IFL	[si] × [si]
Goal	INTENT	INT	[si] × [aco]
	PURPOSE	PRP	[si ∪ ntao] × [si ∪ co ∪ ntao]
Object modifiers	VALUE	VAL	[ql] × [o ∪ si]
	SOURCE	SRC	[loc ∪ ql ∪ ntao ∪ ico] × [o]
Typical syntactic subjects	AGENT	AGT	[aco] × [si]
	EXPERIENCER	EXP	[o] × [si]
	INSTRUMENT	INS	[co ∪ ntao] × [si]
Typical direct objects	THEME	THM	[o] × [ev]
	TOPIC	TPC	[o ∪ si] × [ev]
	STIMULUS	STI	[o] × [ev]
Association	ASSOCIATION	ASO	[ent] × [ent]
	KINSHIP	KIN	[aco] × [aco]
	IS-A	ISA	[ent] × [ent]
	PART-WHOLE	PW	[o] × [o] ∪ [loc] × [loc]
	MAKE	MAK	[co ∪ ntao] × [co ∪ ntao]
	POSSESSION	POS	[co] × [co]
	MANNER	MNR	[ql ∪ st ∪ ntao] × [si]
	RECIPIENT	RCP	[co] × [ev]
	SYNONYMY	SYN	[ent] × [ent]
	AT-LOCATION	AT-L	[o ∪ si] × [loc]
	AT-TIME	AT-T	[o ∪ si] × [tmp]
	PROPERTY	PRO	[ntao] × [o ∪ si]
	QUANTIFICATION	QNT	[qn] × [o ∪ si]

have intentions either, their agents or experiencers might. For example, *The [professor]$^{y}_{aco}$'s goal is to [teach students all the material]$^{x}_{ev}$*; INT(*teach students all the material*, *professor*).

- PURPOSE(*x*, *y*) can be defined for situations, concrete objects, and nontemporal abstract objects, it is somehow a broader relation than INTENT. For example, *Half of the [garage]$^{y}_{ico}$ is used for [storage]$^{x}_{ntao}$*; PRP(*storage*, *garage*).

The object modifiers cluster includes relations that describe **attributes** of objects and situations:

- SOURCE(*x*, *y*) holds if *x* expresses the origin of *y*. *x* could be either a physical location or a mental, information or material origin. For example, *We had a great time with the [Mexican]$^{x}_{ql}$ [students]$^{y}_{aco}$*; SRC(*Mexican*, *students*).
- VALUE(*x*, *y*) holds otherwise. For example, *Not all [smart]$^{x}_{ql}$ [kids]$^{y}_{aco}$ get good grades*; VAL(*smart*, *kids*).

The typical syntactic subjects cluster includes relations that encode links between a **typical syntactic subject** and a situation. The differences rely on the characteristics of the subject and the connection per se:

- AGENT(x, y). x must be volitional, and therefore only animate concrete objects can be part of the domain of AGT. For example, $[\text{John}]^{x}_{aco}$ $[\text{got married}]^{y}_{ev}$ last Spring; AGT(*John*, *got married*).
- EXPERIENCER(x, y). x does not change the situation, only experiences. x does not participate intentionally in y either. The difference between AGENT and EXPERIENCER can sometimes be revealed by the nature of the event. For example, verbs such as *drown*, *find*, and *occur* need an EXPERIENCER, and verbs such as *dive*, *search*, and *think about* require an AGENT. *His [cell phone]*$^{x}_{ico}$ *[suffered]*$^{y}_{ev}$ *some water damage at the pool party*; EXP(*cell phone*, *suffered*).
- INSTRUMENT(x, y). x is used to perform y. x is a tool or device that facilitates y. For example, *[The hammer]*$^{x}_{ico}$ *[broke]*$^{y}_{ev}$ *the window*; INS(*the hammer*, *broke*).

The typical direct objects cluster includes relations encoding **typical syntactic direct objects**: All of them encode attributes of an event since the presence of any of these relations imply a change.

- THEME(x, y) holds if x is affected or directly involved by y, y affects x somehow. For example, *John [read]*$^{y}_{ev}$ *[the book]*$^{x}_{ico}$ *twice*; THM(*the book*, *read*).
- TOPIC(x, y) holds if y is a communication verb, such as *talk* and *argue*. For example, *John [discussed]*$^{y}_{ev}$ *[the issue]*$^{x}_{ntao}$ *too late*; TPC(*the issue*, *discussed*).
- STIMULUS(x, y) holds if y is a perception verb and x and stimulus that makes y happen. y makes x happen somehow. For example, *John [perceived]*$^{y}_{ev}$ *[the ship]*$^{x}_{ico}$ *coming over the horizon*; STI(*the ship*, *perceived*).

The association cluster includes ASSOCIATION and KINSHIP. In the cluster, more specific relations are preferred, if they hold and are more suitable to the situation.

- ASSOCIATION(x, y) is a very broad relation between any pair of entities. For example, *[John]*$^{x}_{aco}$ *and [Mary]*$^{y}_{aco}$ *work at the same company*; ASO(*John*, *Mary*).
- KINSHIP(x, y) encodes a particular relation between relatives. If KINSHIP(x, y), then ASSOCIATION(x, y) holds. For example, *[John]*$^{x}_{aco}$ *visited [his parents]*$^{y}_{aco}$ *for Christmas*; KIN(*John*, *parents*).

The rest of the relations do not fall into any particular cluster. These relations are specific and have their own unique characteristics.

- IS-A(x, y) holds if x is a kind of y.
- PART-WHOLE(x, y) holds if x is part of y.
- MAKE(x, y) holds if x makes or produces y. For example, MAK(*BMW*, *cars*).
- POSSESSION(x, y) holds if y owns x. For example, *[John]*$^{y}_{aco}$*'s [truck]*$^{x}_{ico}$ encodes POS(*truck*, *John*).
- MANNER(x, y) encodes the way in which a situation occurs. For example, MNR(*quick*, *delivery*).

- RECIPIENT(x, y) captures the connection between an event and an object which is the receiver of the event. For example, *John [gave]$^{y}_{ev}$ [Mary]$^{x}_{aco}$ roses* and *John [stole]$^{y}_{ev}$ [Mary]$^{x}_{aco}$'s car.*
- SYNONYMY(x, y) holds if x is a synonym of y.
- AT-LOCATION(x, y) holds if x is at location y. The LOCATION relation denoted by AT-LOCATION in our sample set of twenty-six relations.
- AT-TIME(x, y) holds if x is at time y.
- PROPERTY(x, y) describes links between a situation or object and its characteristics. For example, PRO(*height*, *John*).
- QUANTIFICATION(x, y) holds if y is quantitatively determined by x. For example QNT(*a dozen*, *eggs*).

4 Composition of semantic relations

The goal of composing semantic relations is to acquire new instances of semantic relations by instantiating inference axioms over already identified instances of relations. An axiom takes two semantic relations as input, called premises, and yields a third relation as conclusion. The conclusion reveals implicitly stated knowledge, building an extra layer of semantics which was neglected before. For example, given *"John is in the master bedroom of his condo"*, assume that a semantic parser detects, among others, the following semantic relations: AT-LOCATION(*John*, *master bedroom*) and PART-WHOLE(*master bedroom*, *his condo*), but no connection between *John* and *his condo* is explicitly stated. To connect *John* to *his condo*, we compose AT-LOCATION(*John*, *in master bedroom*) and PART-WHOLE(*master bedroom*, *his condo*) relations. As a result, we conclude that *John* is actually in *his condo* with a valid relation AT-LOCATION(*John*, *in his condo*). In this example, we actually used an axiom which states that *concepts are also located at the whole of its location*, in order to infer the missing relation, AT-LOCATION(*John*, *his condo*).

4.1 Formal definition

We define an axiom by using the composition operator '$\circ$'. Formally, $R_1 \circ R_2 \rightarrow R_3$, where R_1 and R_2 are the premises and R_3 the conclusion. The composition function has several properties.

Relation-type compatibility: Two relations R_1 and R_2 are compatible iff $\text{RANGE}(R_1) \cap \text{DOMAIN}(R_2) \neq \emptyset$. Let us denote $\text{RANGE}(R_1) \cap \text{DOMAIN}(R_2) = I$. Unless $I = \text{RANGE}(R_1) = \text{DOMAIN}(R_2)$, a restriction takes place when combining the two relations. A *backward restriction* takes place if $\text{RANGE}(R_1) \neq I$ and a *forward restriction* if $\text{DOMAIN}(R_2) \neq I$. In the former case, $\text{RANGE}(R_1)$ is reduced; in the later, $\text{DOMAIN}(R_2)$ is reduced. For example, in $\text{THM}^{-1}(x, y)$ o $\text{AT-L}(y, z)$; y is forward restricted to objects [ev]×[o] ; [o ∪ si]×[loc]. Another example for backward restriction is $\text{AT-L}^{-1}(x, y)$ o $\text{THM}(y, z)$; y is backward restricted to events as [loc]×[o ∪ si] ; [o]×[ev]. Note that one can find a forward and backward restriction with the same pair of relations.

Conclusion closure: The composition of $R_1 \circ R_2 \rightarrow R_3$ is said to be closed under composition, if there exists an R_3 in the set of relations where $\text{RANGE}(R_3) \subseteq$

RANGE(R$_2$) and DOMAIN(R$_3$) $\subseteq$ DOMAIN(R$_1$). There are restrictions for DOMAIN and RANGE. The RANGE restriction exists when RANGE(R$_3$) $\subset$ RANGE(R$_2$) and the DOMAIN restriction exists when DOMAIN(R$_3$) $\subset$ DOMAIN(R$_2$).

Conclusion validity: In $R_1(x, y) \circ R_2(y, z) \rightarrow R_3(x, z)$, R_3 is semantically valid if it holds with arguments x and z. In the example of *"John is in the master bedroom of his condo"*, the fact that master bedroom is a part of *"John's condo"* semantically suggests that any concept that is in the bedroom is also in the condo, making AT-LOCATION(x, y) $\circ$ PART-WHOLE(y, z) $\rightarrow$ AT-LOCATION(x, z) composition a valid axiom.

Dominance: If R_3 is one of R_1 or R_2 then, we talk about the dominance of any premises. R_1 **left dominates** R_2, denoted by $R_1 \triangleright R_2$, iff the composition of R_1 and R_2 yields R_1. Formally, $R_1 \triangleright R_2$ iff $R_1(x, y) \circ R_2(y, z) \rightarrow R_1(x, z)$. R_1 **right dominates** R_2, denoted by $R_2 \triangleleft R_1$, iff the composition of R_2 and R_1 yields R_1. Formally, $R_2 \triangleleft R_1$ iff $R_2(x, y) \circ R_1(y, z) \rightarrow R_1(x, z)$. R_1 **completely dominates** R_2, denoted by $R_1 \bowtie R_2$, iff $R_1 \triangleright R_2$ and $R_2 \triangleleft R_1$. Formally, $R_1 \bowtie R_2$ iff $R_1(x, y) \circ R_2(y, z) \rightarrow R_1(x, z)$ and $R_2(x, y) \circ R_1(y, z) \rightarrow R_1(x, z)$. In general, if one only states that R_1 *dominates* R_2, it should be interpreted as R_1 *completely dominates* R_2.

In order to apply the composition function over R_1 and R_2 they must fulfill the following necessary conditions: (i) R_1 and R_2 must be type compatible (relation-type compatibility); (ii) R_3 must be closed over the set of relations being used (conclusion closure); (iii) R_3 must be a valid conclusion which holds with its premises (conclusion validity).

In principle, one could think of applying the composition for every single pair of relations. However, using the type compatibility, one can easily filter the entire spectrum of pair wise combinations to a smaller set of compatible pairs. However, there are still compatible pairs that do not result in a closed conclusion. For example, in QUANTIFICATION $\circ$ AT-LOCATION $\rightarrow$ R_3, QUANTIFICATION (QNT), and AT-LOCATION (AT-L) are type compatible relations, but they do not compose a closed relation, since there is no R_3 in the relation set, where DOMAIN(R$_3$) $\subseteq$ DOMAIN(QNT) and RANGE(R$_3$) $\subseteq$ RANGE(AT-L). Therefore, we can further filter those pairs that do not yield a closed conclusion.

It is important to note that a pair of compatible relations that yields a closed conclusion relation may not produce a valid conclusion. We call the composition that produces a valid conclusion a valid axiom. Therefore, not any given pair of compatible relations forms a valid axiom. For example, KINSHIP and AT-LOCATION are compatible and closed, but yield no valid inference. Given KIN(*Mary*, *John*) and LOC(*John*, *Dallas*), no relation can be inferred between *Mary* and *Dallas*. Thus, the composition is not an axiom.

An axiom $R_1(x, y) \circ R_2(y, z) \rightarrow R_3(x, z)$ is instantiated by the instances of premises $R_1(x, y)$ and $R_2(y, z)$ to generate a conclusion $R_3(x, z)$.

As a general rule, any instantiation of an axiom holds. However, sometimes only certain instantiations in specific contexts hold. We name such an axiom PLAUSIBLE axiom and denote as $R_1 \circ R_2 \rightarrow$ PLAUSIBLE R_3.

In this study, we are interested in axioms that yield strong and meaningful inferences and thus we do not count PLAUSIBLE axioms in the results. The goal is to find axioms that work regardless of particular instantiations of relations. At the same time, the aim is to extract concrete and narrow knowledge. In this paper, we focus on the composition of AT-LOCATION and other relations.

If no relation from the given set holds between two concepts, then the relation called OTHER holds. OTHER is represented by the symbol $\perp$. Formally, $\perp(x, y)$ iff $\neg\exists R_i$ such that $R_i(x, y)$.

We define the inverse representation of a relation R, as $R^{-1}(y, x)=R(x, y)$, in short R^{-1}, and R^{-1}inherits all properties of R. The inverse representation of an axiom is, then, defined as $(R_i \circ R_j)^{-1}=R_j^{-1} \circ R_i^{-1}$.

Lemma 1: $R_i \circ R_j = (R_j^{-1} \circ R_i^{-1})^{-1}$. The composition of R_i and R_j is equal to the inverse composition of R_j^{-1} and R_i^{-1}.

Lemma 2: $\perp^{-1} = \perp$. If there is no relation from one concept to another, then there is also no reverse representation. $R(x, y) = R^{-1}(y, x)$; therefore, $\perp^{-1}(y, x)$ iff $\neg\exists R_i^{-1}$ such that $R_i^{-1}(y, x)$.

Lemma 3: $\forall i : \perp \bowtie R_i$. The $\perp$ completely dominates any relation. The $\perp$ left dominates any relation; $\forall i : \perp(x, y) \circ R_i(y, z) \rightarrow \perp(x, z)$, then $\perp \rhd R_i$. Similarly, the $\perp$ right dominates any relation; $\forall i : R_i(x, y) \circ \perp(y, z) \rightarrow \perp(x, z)$, then $\perp \lhd R_i$. Therefore, $\forall i : R_i \bowtie \perp$.

Using the notation above, we can state the following properties:

- R is transitive iff $R(x, y) \circ R(y, z) \rightarrow R(x, z)$.
- R is reflexive iff $\forall$ x: $R(x, x)$.
- R is symmetric iff $\forall$ x,y: there exist $R(x, y)$ and $R(y, x)$.

4.2 Unique composition of two relations

For a given pair of relation types, $R_1(x, y)$ and $R_2(y, z)$, there are four unique combinations for composition, shown in Table 2.

Any one of these combinations is subject to type compatibility, closure, and validity to qualify for becoming a valid axiom. We classify a pair as a compatible-pair if any one of four combinations is compatible. Similarly, we classify a pair as a closed-pair if any of its four combinations is closed under the relation set that we use.

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 R_i and R_j and their inverses R_i^{-1} and R_j^{-1}, there are $4 \times 4 = 16$ different possible combinations. Due to the property of $R_i \circ R_j = (R_j^{-1} \circ R_i^{-1})^{-1}$, there are $\binom{4}{2} = 6$ repeating compositions in this set of 16. The remaining 10 compositions are unique: (i) 4 combine R_1, R_2 and their inverses; (ii) 3 combine R_1 and its inverse; (iii) 3 combine R_2 and its inverse.

Table 2. *The four unique axioms taking* R$_1$ *and* R$_2$ *as premises and yielding a valid conclusion relation* R$_3$ *with a diagonal arrow*

R$_1$ ∘ R$_2$	R$_1$ ∘ R$_2$$^{-1}$	R$_2$ ∘ R$_1$	R$_2$$^{-1}$ ∘ R$_1$
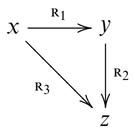 case 1	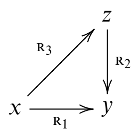 case 2	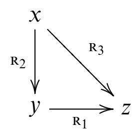 case 3	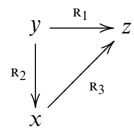 case 4

Therefore, for n relations there are $2n^2 + n$ potential axioms as given below:

$$\binom{n}{2} \times 4 + 3n = 2 \times n(n-1) + 3n = 2n^2 - 2n + 3n = 2n^2 + n$$

When working with twenty-six relations, there are 1,378 potential unique axioms. Many can be discarded by the compatibility test; some compatible pairs do not yield any valid inference.

5 Composition of AT-L relation with other relations

In this section, we analyze the composition of AT-L with other relations in the previously explained set of twenty-six relations (see Table 3). The analysis is done on a cluster basis.

5.1 AT-L relation and reason *cluster*

If any of the arguments of CAU, JST, or IFL relations is at a location, then the other argument is also at the same location (see case 1 and case 4 in Table 4). The case 2 and case 3 are not type compatible for composition. The conclusion is not valid, if both situations are having different locations that are explicitly stated.

Examples:

- *"Since John drove his car carelessly in the city, he had an accident."* The following axiom infers that the accident also happened in the city.
  AT-L^{-1}(*in the city*, *drove carelessly*) ∘ CAU(*drove carelessly*, *accident*) → AT-L^{-1}(*in the city*, *accident*).
- *"John drove his car carelessly. Because of that he had an accident on the street."* CAU(*drove carelessly*, *accident*) ∘ AT-L(*accident*, *on the street*) → AT-L(*drove carelessly*, *on the street*). The axiom suggests that the act of driving carelessly also happened on the street.
- *"They do not smoke in the restaurant, because it is forbidden."* The axiom can infer that smoking is forbidden in the restaurant. JST(*forbidden*, *do not smoke*) ∘ AT-L(*do not smoke*, *in the restaurant*) → AT-L(*forbidden*, *in the restaurant*).

Table 3. *Semantic composition of twenty-six relations with* AT-LOCATION *relation*

			AT-L		
Cluster	Relation	Domain × Range	Compatibility	Closure	Validity
Reason	CAU	[si] × [si]	Yes	Yes	AT-L
	JST	[si ∪ ntao] × [si]	Yes	Yes	AT-L
	IFL	[si] × [si]	Yes	Yes	AT-L
Goal	INT	[si] × [aco]	Yes	Yes	AT-L
	PRP	[si∪ntao] × [si∪co∪ntao]	Yes	Yes	AT-L
Object modifiers	VAL	[ql] × [o ∪ si]	Yes	No	-
	SRC	[loc∪ql∪ntao∪ico] × [o]	Yes	Yes	-
Typical syntactic subjects	AGT	[aco] × [si]	Yes	Yes	AT-L
	EXP	[o] × [si]	Yes	Yes	AT-L
	INS	[co ∪ ntao] × [si]	Yes	Yes	AT-L
Typical direct objects	THM	[o] × [ev]	Yes	Yes	AT-L
	TPC	[o ∪ si] × [ev]	Yes	Yes	-
	STI	[o] × [ev]	Yes	Yes	-
Association	ASO	[ent] × [ent]	Yes	Yes	-
	KIN	[aco] × [aco]	Yes	Yes	-
	ISA	[ent] × [ent]	Yes	Yes	AT-L
	PW	[o] × [o] ∪ [l] × [l] ∪ [t] × [t]	Yes	Yes	AT-L
	MAK	[co ∪ ntao] × [co ∪ ntao]	Yes	Yes	AT-L
	POS	[co] × [co]	Yes	Yes	-
	MNR	[ql ∪ st ∪ ntao] × [si]	Yes	Yes	-
	RCP	[co] × [ev]	Yes	Yes	-
	SYN	[ent] × [ent]	Yes	Yes	AT-L
	AT-T	[o ∪ si] × [tmp]	Yes	No	-
	PRO	[ntao] × [o ∪ si]	Yes	Yes	-
	QNT	[qn] × [o ∪ si]	Yes	No	-

- *"They do not smoke, because it is forbidden in the restaurant."* The axiom can infer that no-smoking happens in the restaurant. AT-L^{-1}(*in the restaurant, forbidden*) ∘ JST(*forbidden, do not smoke*) → AT-L^{-1}(*in the restaurant, do not smoke*).
- *"John misses classes at UTD; therefore his grades are low."* The axiom can infer that John's grades are also at UTD. AT-L^{-1}(*UTD, missing classes*) ∘ IFL(*missing classes, low grades*) → AT-L^{-1}(*UTD, low grades*).

Since all relations (CAU, JST, IFL) in the reason cluster are close in meaning and give the same valid conclusion, we treat this composition between AT-L and reason cluster, as seen in Table 4.

Table 4. *The four axioms for* AT-L *and* reason *(N/C stands for not compatible)*

reason ∘ AT-L	reason ∘ AT-L^{-1}	AT-L ∘ reason	AT-L^{-1} ∘ reason

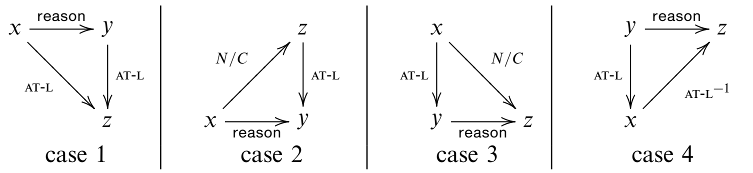

Table 5. *The four axioms for* AT-L *and* goal *(N/C stands for not compatible and N/V stands for not valid)*

goal ∘ AT-L	goal ∘ AT-L^{-1}	AT-L ∘ goal	AT-L^{-1} ∘ goal

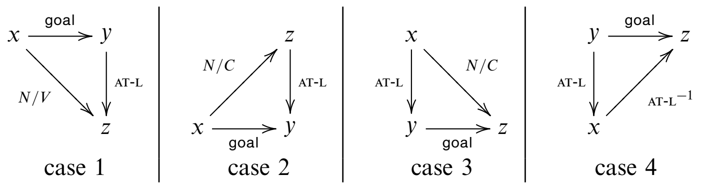

5.2 AT-L relation and goal *cluster*

If the intention INT or purpose PRP takes place at a location, then the concept which has the intention or purpose shares the same location (see case 4 in Table 5).
Examples:

- *"Professor intends to teach"*; INT(*teach*, *professor*). The professor may teach in the future. If the professor is at a location now or any time in the future, there is no relation between the location of the professor and the intention. However, if there is a location for the intention such as AT-L(*teach*, *in classroom*), then the professor has an intention to be in the classroom and probably he will be. AT-L^{-1}(*in classroom*, *teach*) ∘ INT(*teach*, *professor*) → AT-L^{-1}(*in classroom*, *professor*).
- *"The purpose for a blue dress is to look pretty at the party."* AT-L^{-1}(*at the party*, *look pretty*) ∘ PRP(*look pretty*, *blue dress*) → AT-L^{-1}(*at the party*, *blue dress*).

Since both relations (INT and PRP) in goal cluster are close in meaning and give the same valid conclusion, we treat this composition between AT-L and goal.

The goal ∘ AT-L composition (see case 1 in Table 5) is not valid. Since intentions occur in the future, we cannot infer that the concept and its intention should be at the same location because animate objects may change their locations with time.

5.3 AT-L relation and object modifier *cluster*

The VAL relation is type compatible with AT-L in case 1; however, it is not closed under the relation set (see also Table 3). The SRC relation does not result in a valid conclusion, since source of a concept may not share a location.

Table 6. *The four axioms for* AT-L *and* syn-sub *(N/C stands for not compatible)*

syn-sub ∘ AT-L	syn-sub ∘ AT-L^{-1}	AT-L ∘ syn-sub	AT-L^{-1} ∘ syn-sub

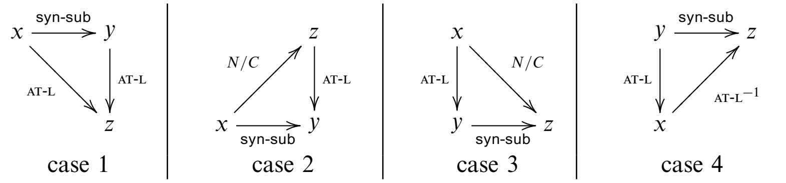

5.4 AT-L relation and **typical syntactic subjects** *cluster*

If the animate object x is the agent (AGT) for a situation y and the situation y is at a location z, then the agent is also at that location z (see case 1 in Table 6). The reverse is also valid as in case 4 of Table 6, meaning that if the agent is at a location, it means that the situation(event) is also at that location.

Examples:

- *"John got married in Dallas"* is encoded by following two relations: AGT(*John*, *married*) and AT-L(*married*, *Dallas*). Therefore, John was in Dallas at some point, AT-L(*John*, *Dallas*).
  AGT(*John*, *married*) ∘ AT-L(*married*, *Dallas*) → AT-L(*John*, *Dallas*).
- *"While in Dallas, John got married"*; AGT(*John*, *married*), AT-L(*John*, *Dallas*). These statements conclude that AT-L(*married*, *Dallas*).
  AT-L^{-1}(*Dallas*, *John*) ∘ AGT(*John*, *married*) → AT-L^{-1}(*Dallas*, *married*).

EXP relation follows a similar composition with AT-L. If the situation being experienced (EXP) is at a location, then the experiencer is (at some point in time) at the same location.

Examples:

- *"The cell phone suffered water damage in the swimming pool."*
  EXP(*cell-phone*, *suffered*) ∘ AT-L(*suffered*, *in the swimming-pool*) → AT-L(*cell-phone*, *in the swimming-pool*).
- The reverse is similarly true. The fact that *"the cell-phone is in the swimming pool"* and *"the cell-phone experiences suffering"* conclude that suffering also happens in the swimming pool.
  AT-L^{-1}(*swimming pool*, *cell-phone*) ∘ EXP(*cell-phone*, *suffering*) → AT-L^{-1}(*swimming pool*, *suffering*).

The INS has the same conclusion in the cluster. If the situation y, which is facilitated by the instrument (or tool) *x*, is at a location, then the instrument x must be at that location at some point in time.

Examples:

- *"The scissors cut the paper in the house"* is encoded by these two relations: INS(*scissors*, *cutting*) and AT-L(*cutting*, *in the house*). These two relations can

Table 7. *The four axioms for* AT-L *and* THM *(N/C stands for not compatible)*

THM ∘ AT-L	THM ∘ AT-L^{-1}	AT-L ∘ THM	AT-L^{-1} ∘ THM

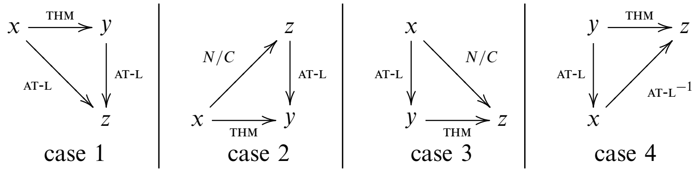

give the following axiom stating the location for the tool of *"scissors"*.
INS(*scissors*, *cutting*) ∘ AT-L(*cutting*, *in the house*) → AT-L(*scissors*, *in the house*)

- The fact that scissors are in the house may mean that the cutting happens in the house.
AT-L^{-1}(*house*, *scissors*) ∘ INS(*scissors*, *cutting*) → AT-L^{-1}(*house*, *cutting*)

Since all relations (AGT, EXP, and INS) in typical syntactic subjects (syn-sub)cluster are close to each other in meaning and give the same valid conclusion, we treat this composition between AT-L and syn-sub cluster as shown in Table 6.

5.5 AT-L relation and typical direct objects *cluster*

The relations (THM, TPC, and STI) in typical direct objects (drct-obj) cluster behave differently. For the composition of THM(*x*, *y*) and AT-L(*y*, *z*), if the event y takes place at a certain location z, then the location for the theme *x* is also *z*. Table 7 plots the four potential compositions of the pair.

Examples:

- The statement *"John drew a graph in his room"* will result in two relations: THM(*graph*, *drew*) and AT-L(*drew*, *in his room*). Composing these two relations in this order will give a valid axiom:
THM(*graph*, *drew*) ∘ AT-L(*drew*, *in his room*) → AT-L(*graph*, *in his room*)
- The location of the theme could also infer the location of the activity that is associated with the theme. Following the same example above, if the graph is in the room, then the drawing probably happened in the room as well. Therefore, we instantiate another axiom below:
AT-L^{-1}(*in his room*, *graph*) ∘ THM(*graph*, *drew*) → AT-L^{-1}(*in his room*, *drew*)

For TPC, the conclusion is not valid. The event and the topic of the event do not share a location.

Example:

- In the statement *"The topic of the talk is about cars."* The talk takes place in the classroom, but the cars are not in the class room. Therefore, the composition is not an axiom.
TPC(*cars*, *talk*) ∘ AT-L(*talk*, *in the classroom*) ↛ AT-L(*cars*, *in the classroom*)

Table 8. *The four axioms for* AT-L *and* ISA

ISA ∘ AT-L	ISA ∘ AT-L^{-1}	AT-L ∘ ISA	AT-L^{-1} ∘ ISA

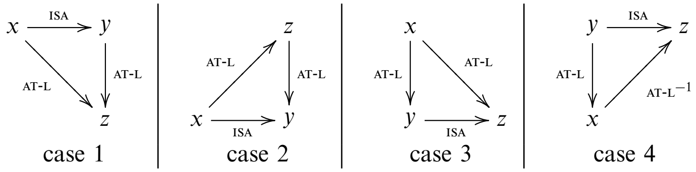

Similarly, the stimulant (STI) and the event that is stimulated may not share a location, since they may be at a quite some distance.

Example:

- In the statement *"John heard the train on the beach"*, the fact that *"John is on the beach"* does not mean anything about the location of the train.
  STI(*train*, *hearing*) ∘ AT-L(*hearing*, *on the beach*) ↛ AT-L(*train*, *on the beach*)

5.6 AT-L relation and **Association** *cluster*

In ASO, being associated for a pair of concepts does not carry any locative feature. Similarly, any two concepts that are related to kinship, KIN, do not share a location.

5.7 AT-L relation and other unclustered relations

The full composition of ISA relation with AT-L produces four valid axioms (see Table 8). Simply, both concepts of an ISA relation share the same location because of its inheritance feature.

Example:

- In a text *"Rice in Houston is a research university"*, we would have AT-L(*Rice*, *Houston*) and ISA(*Rice*, *research university*), which can be composed by the axiom in case 4 resulting below valid AT-L instance.
  AT-L^{-1}(*Houston*, *Rice*) ∘ ISA(*Rice*, *research university*) → AT-L^{-1}(*Houston*, *research university*)

Axioms involving PW are studied under three subtypes of PW, as defined in WordNet (Miller 1995): (i) Part, PW_p; (ii) Substance, PW_s; and (iii) Member PW_m. Most examples are from WordNet, especially the PART-WHOLE annotation and definitions. Table 9 summarizes the aggregated axioms for AT-L and PW semantic composition for all subtypes. The table shows a valid axiom for a case, if at least one subtype in that case produces a valid conclusion.

Case 1: PW(x, y) ∘ AT-L(y, z)

This axiom is relatively straightforward and it captures the fact that the parts, substances and members of a whole are at the same location as the whole.

Table 9. *The four axioms for* AT-L *and* PW *(N/V stands for not valid)*

PW ∘ AT-L	PW ∘ AT-L^{-1}	AT-L ∘ PW	AT-L^{-1} ∘ PW

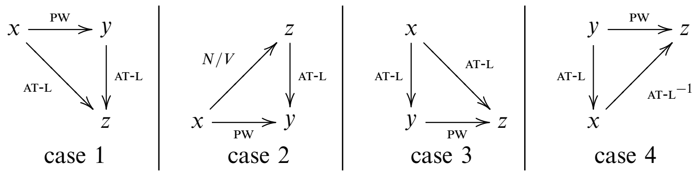

(a) PW$_p(x, y)$ ∘ AT-L$(y, z) \rightarrow$ AT-L(x, z)

Example:

- Given PW$_p$(*horn button*, *car horn*) and AT-L(*car horn*, *garage*), then AT-L(*horn button*, *garage*).

(b) PW$_s(x, y)$ ∘ AT-L$(y, z) \rightarrow$ AT-L(x, z)

Example:

- Given PW$_s$(*caffeine*, *coffee*) and AT-L(*coffee*, *pot*), then AT-L(*caffeine*, *pot*).

(c) PW$_m(x, y)$ ∘ AT-L$(y, z) \rightarrow$ AT-L(x, z)

Example:

- Given PW$_m$(*Shiite*, *Shiah*) and AT-L(*Shiah*, *Iran*), then AT-L(*Shiite*, *Iran*).

Case 2: PW(x, y) ∘ AT-L$^{-1}(y, z)$

The location y of an object z is the whole of a part x. The case analyzes the potential axioms for all types of PW.

(a) PW$_p(x, y)$ ∘ AT-L$^{-1}(y, z) \nrightarrow$ AT-L(x, z)

For PW$_p$, the fact that something is located at a whole does not necessarily mean that it is also located at a particular part of the whole.

Example:

- Maquiladora is a plant in Mexico AT-L(*Maquiladora*, *Mexico*) and Acapulco is a part of Mexico as stated in WordNet PW$_p$(*Acapulco*, *Mexico*). These two relations cannot conclude that the plant is actually in Acapulco as it may be in any part of Mexico. Therefore, we conclude that the axiom for this case does not hold.

(b) PW$_s(x, y)$ ∘ AT-L$^{-1}(y, z) \nrightarrow$ AT-L(x, z)

For substance, the conclusion is also similar and does not hold. If a whole constitutes a location for an object or situation, we cannot conclude that the substance of the whole constitutes a location.

(c) PW$_m(x, y)$ ∘ AT-L$^{-1}(y, z) \nrightarrow$ AT-L(x, z)

The member-collection Part-Whole also does not hold. If the collection is a location for an object or situation, then there is no reasonable indication for the member to constitute a location.

Case 3: AT-L(x, y) ∘ PW(y, z)

(a) AT-L(x, y) ∘ PW$_p(y, z) \rightarrow$ AT-L(x, z)

If the location *y* of a concept *x* is part of a whole *z*, then *x* is located in *z*.

Example:

- In the statement *"Myalgia is a pain in a muscle"*, AT-L(*myalgia*, *muscle*); and *"muscle is one of the contractile organs of the body"*, PW_p(*muscle*, *body*). Therefore, AT-L(*myalgia*, *body*).

(b) AT-L(x, y) $\circ$ $PW_s(y, z) \rightarrow$ AT-L(x, z)
If the location *y* of a concept *x* is a substance of *z*, then *x* is located in *z*.
In regular text, one rarely states a location which is a substance of another concept. However, from a theoretical point of view, the axiom holds.

Example:

- In the statement *exostosis* is defined as *"a benign outgrown located on bone"*, AT-L(*exostosis*, *bone*). If one finds exostosis on the particular bone that is a substance of a horn, PW_s(*bone*, *horn*), then one can infer that AT-L(*exostosis*, *horn*).

(c) AT-L(x, y) $\circ$ $PW_m(y, z) \rightarrow$ PLAUSIBLE AT-L(x, z)
If the location *y* of a concept *x* is a member of a whole *z*, then *x* might be located in *z*. In this case, the conclusion is a plausible relation, since depending on the nature of the PW_m the inference might not be valid. We note that the sorts of concepts involved in the PW_m are the key. If they belong to different hierarchies in the ontology of sorts, the inference does not hold.

Example:

- If *"a [meeting] occurs at a [restaurant]$_{co}$ which is a member of a particular [restaurant chain]$_{ao}$"*, one cannot clearly conclude that the meeting occurs in the restaurant chain. On the other hand, if a [bug] is in the oldest [tree]$_{co}$ belonging to the Sequoia National [Forest]$_{co}$, one can state that the bug is in the Sequoia National Forest.

Case 4: AT-L$^{-1}(x, y)$ $\circ$ PW(y, z)

(a) AT-L$^{-1}(x, y)$ $\circ$ $PW_p(y, z) \rightarrow$ PLAUSIBLE AT-L$^{-1}(x, z)$
The fact that a part *y* of a whole *z* is at a location *x* does not necessarily mean that *z* is also located at *x*. As a general rule, if the whole *z* is not larger than the location *x*, AT-L(x, z) holds.

Positive example:

- WordNet states PW_p(*time-ball*, *timepiece*). If the particular *time-ball* which is part of a *timepiece* is in an observatory, AT-L^{-1}(*observatory*, *time-ball*), then AT-L(*timepiece*, *observatory*).

Negative example:

- As a counter example, given AT-L^{-1}(*Panhandle*, *Lubbock*) and PW_p(*Lubbock*, *Texas*), one cannot infer that AT-L^{-1}(*Panhandle*, *Texas*), which is equal to AT-L(*Texas*, *Panhandle*). In this particular case, the whole *Texas* is bigger than the location *Panhandle*; therefore, the inference does not hold. With the existence of positive and negative examples, we conclude that the axiom generates a plausible relation.

Table 10. *Summary of Axioms for* PART-WHOLE *in three subtypes: (1)* PW_p*: part, (2)* PW_s*: substance, and (3)* PW_m*: member*

Case	Premises		Conclusion
3	AT-L	PW_p PW_s	AT-L
		PW_m	PLAUSIBLE AT-L
4	AT-L^{-1}	PW_p PW_s PW_m	PLAUSIBLE AT-L^{-1} PLAUSIBLE AT-L^{-1} PLAUSIBLE AT-L^{-1}
1	PW_p PW_s PW_m	AT-L	AT-L
2	PW_p PW_s PW_m	AT-L^{-1}	NOT-VALID

(b) AT-L$^{-1}(x, y) \circ PW_s(y, z) \rightarrow$ PLAUSIBLE AT-L$^{-1}(x, z)$

$PW_s(x, y)$ implicitly guarantees that x is somehow merged and not easily separable from y. y is made, among others, of x; when observing or dealing with y one sees the whole, and not the different substance it is made of. Therefore, the location of a substance determines the location of the whole.

Example:

- If AT-L^{-1}(*glass*, *tequila*) and PW_P(*tequila*, *margarita*), then AT-L^{-1}(*glass*, *margarita*).

(c) AT-L$^{-1}(x, y) \circ PW_m(y, z) \rightarrow$ PLAUSIBLE AT-L$^{-1}(x, z)$

The fact that a particular member y of a whole z is at a location x does not necessarily mean that z is also located at x. Only if all the members of z are located in x, one can conclude that z is located in x.

Positive example:

- For example, biota is defined as *"the plant and animal life in a particular region"*, and has as members *vegetation* and *fauna*. Therefore, if AT-L^{-1}(*Amazonas*, *vegetation*) and PW_m(*vegetation*, *biota*), then AT-L^{-1}(*Amazonas*, *biota*) holds.

Negative example:

- A counter example is the following. A particular professor is a member of the faculty, PW_m(*professor*, *faculty*). Just because that professor is in *Cancun*, AT-L^{-1}(*Cancun*, *professor*), we cannot conclude that AT-L^{-1}(*Cancun*, *faculty*). The inference would be valid if all the professors were in Cancun.

The summary of all PW-related axioms can be seen in Table 10.

For MAK relation, the composition produces one valid axiom (see Table 11). If a concept y makes or produces another concept z, MAK(y, z) and if asry is at a location

Table 11. *The four axioms for* AT-L *and* MAK *(N/C stands for not compatible and N/V stands for not valid)*

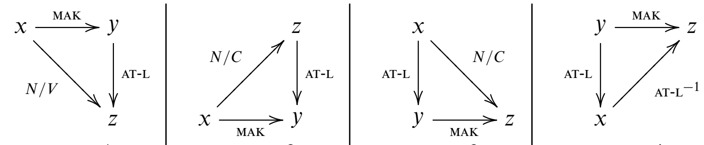

MAK ∘ AT-L	MAK ∘ AT-L^{-1}	AT-L ∘ MAK	AT-L^{-1} ∘ MAK
case 1	case 2	case 3	case 4

x, AT-L(y, x), then z is also at the same location at least some point in time, even though the product may be transferred to some other location later on.

Examples:

- If GM is located in Arlington and producing trucks, then there are trucks in Arlington, as well.
  AT-L^{-1}(*in Arlington*, *GM*) ∘ MAK(*GM*, *trucks*) → AT-L^{-1}(*in Arlington*, *trucks*)

- However, the reverse is not true. If the product is at a location at some point, it does not mean that the producer/maker is also at the same location.
  MAK(*GM*, *trucks*) ∘ AT-L(*trucks*, *in Arlington*) ↛ AT-L(*GM*, *in Arlington*)

For POS relation, none of the compositions produces an axiom. Animate objects (people in this case) and their possessions may not share a location; therefore we cannot conclude for the composition of AT-L and POS.

Examples:

- POS(*Apartment*, *Oscar*) ∘ AT-L(*Oscar*, *in Barcelona*) ↛ AT-L(*Apartment*, *in Barcelona*)
- AT-L^{-1}(*in Barcelona*, *Oscar*) ∘ MAK(*Oscar*, *Apartment*) ↛ AT-L^{-1}(*in Barcelona*, *Apartment*)

For MNR relation, the composition does not produce an axiom. The manner and the location of the situation do not make sense.

For RCP relation, the axiom is only plausible. It holds for some verbs where agent and the event may share a location. For example *"Give"*, *"Receive"* share a location with the agent. However, the events such as *"Steal"* and *"Lose"* may not share a location with the agent; therefore no conclusion can be made.

Positive example:

- *"They give the award to Mary at the banquet."*
  RCP(*Mary*, *give*) ∘ AT-L(*give*, *at the banquet*) → AT-L(*Mary*, *at the banquet*)

Negative example:

- *"Someone stole Mary's car."* Since stealing does not have to happen at the same location as its recipient's, we have the below invalid axiom.

  RCP(*Mary*, *stole*) ∘ AT-L(*stole*, *on the street*) ↛ AT-L(*Mary*, *on the street*)

Table 12. *The four axioms for* AT-L *and* SYN

SYN ∘ AT-L	SYN ∘ AT-L^{-1}	AT-L ∘ SYN	AT-L^{-1} ∘ SYN

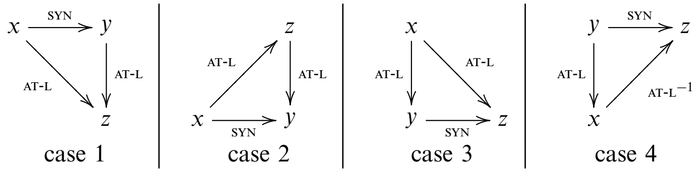

The SYN relation can compose with AT-L relation in all four cases and produce valid conclusions, as shown in Table 12.

Example:

- *"Hershey is a town in Pennsylvania and Keystone State is a synonym for Pennsylvania."*
  AT-L(*Hershey*, *Pennsylvania*) ∘ SYN(*Pennsylvania*, *Keystone State*) → AT-L(*Hershey*, *Keystone State*)

The AT-T relation is not closed under the given set of semantic relations (see Table 3).

PRO does not share a location with the property holder; therefore, the composition does not produce a valid axiom.

The QNT relation is not closed under the given set of semantic relations (see Table 3).

6 Implementation and results

In this section, we report the results of the proposed method for the composition of AT-L relation with other relations. There are three clusters and five relations that *fulfill* the necessary conditions of type compatibility and conclusion closure and produce valid axioms (see Table 3). We apply these axioms to annotated relations in PropBank, WordNet, and eXtended WordNet-Knowledge Base (XWN-KB) to produce new instances of relations. Then, we report the precision of these newly generated instances of relations. We use annotation from PropBank for the reason, typical syntactic subjects, and typical direct obj. clusters along with AT-L relation. Following the proposed method, the implementation first searches for all the sentences in PropBank that have the two relations being composed (e.g., AT-L and THM) by a valid axiom, say $R_1(x, y)$ and $R_2(y, z)$, which share a common concept y. Then it produces the conclusion $R_3(x, z)$.

For annotated and disambiguated PW relations, we use WordNet Meronymy between synsets in three subtypes: (i) Is-Part; (ii) Is-Member; and (iii) Is-Substance, which were previously discussed in Section 5.7.

For all other relations, we use annotated results from XWN-KB. XWN-KB is an upper ontology built as an extension to eXtended WordNet[2] (XWN) (Mihalcea and Moldovan 2001), which is derived from WordNet (WN). The novelty that XWN-KB brings is that the glosses of synsets have been transformed into semantic relations by using a reliable in-house semantic parser, and verified by human annotators. The semantic parser uses a variety of syntactic realizations to generate semantic relations. For example, *knife blade* contains PW(*blade*, *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*). 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 semantic parser 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 whether a relation holds. Four different algorithms are used: decision trees, Naive Bayes, support vector machine, 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). In addition, human-coded rules are used in a small number of unambiguous cases. The semantic parser participated in the SemEval 2007 Task 4 (Badulescu and Srikanth 2007). The result is a knowledge base that is highly interconnected. Unlike a domain-specific ontology that is based on open source information, such as news, intelligence reports, etc., XWN-KB uses definitional glosses of WordNet for its synsets.

6.1 Reason *cluster*

For the reason cluster, both PropBank and XWN-KB annotated relations were available for the application of the valid axioms (see Table 13). In PropBank, there were 46 sentences that have CAU and AT-L relations with a common shared concept, only seven of which made no sense (negative). Below are positive and negative actual sentences.

Positive: *"A half-dozen Soviet space officials , in Tokyo in July for an exhibit, stopped by to see their counterparts at the National Space Development Agency of Japan."* In this example, we have below valid axiom.

CAU(*exhibit*, *stop*) ∘ AT-L(*stop*, *Japan*) → AT-L(*exhibit*, *Japan*)

Negative: *"Here, at a soccer stadium near the black township of Soweto yesterday, were eight leaders of the African National Congress, seven of whom had spent most of their adult lives in prison for sabotage and conspiracy to overthrow the government."* This example does not work for the compositional axiom below.

CAU(*sabotage and conspiracy*, *spending*) ∘ AT-L(*spending*, *in prison*) ↛ AT-L(*sabotage and conspiracy*, *in prison*)

[2] http://xwn.hlt.utdallas.edu/

Table 13. *Results for LOCATION axioms*

		AT-L			
		PropBank		WordNet + XWN-KB	
Cluster	Relation	Instances	Precision	Instances	Precision
Reason	CAU JST IFL	46	39/46 = 85%	38	38/38 = 100%
Goal	INT PRP	N/A	N/A	91	88/91 = 97%
Typical syntactic subjects	AGT EXP INS	4053	91/100 = 91%	503	100/100 = 100%
Typical direct object	THM	138	97/100 = 97%	832	96/100 = 96%
	ISA	N/A	N/A	1997	100/100 = 100%
	PW	N/A	N/A	179	100/100 = 100%
	MAK	N/A	N/A	8	8/8 = 100%
	SYN	N/A	N/A	264	98/100 = 98%

In this example, spending and the causing concept do not share the same location.

For PropBank, the precision achieved is 39/46 = 85 per cent. We then repeat the same exercise for WordNet and XWN-KB, and we find thirty-eight instances where AT-L and reason cluster relations share a common concept in a composition. We checked all thirty-eight and found that all thirty-eight make sense, resulting in hundred per cent precision.

6.2 **Goal** *cluster*

For goal cluster, only XWN-KB annotated relations were available. There were only ninety-one cases where both AT-L and a goal cluster relation (INT or PRP) share a common concept. We checked all ninety-one and eighty-eight made sense giving a precision of ninety-seven per cent.

6.3 **Typical syntactic subjects** *cluster*

There are 4,053 sentences in PropBank that have AGENT and AT-LOCATION relations sharing a common concept. We randomly selected one hundred of these, and only nine made no sense (negative), resulting ninety-one per cent precision. An example sentence is given below:

"In Seoul, officials began visiting about 26,000 cigarette stalls to remove illegal posters and signboards advertising imported cigarettes". The axiom that produces a new AT-LOCATION relation is given below.

AGT(*officials*, *began*) $\circ$ AT-L(*began*, *in Seoul*) $\rightarrow$ AT-L(*officials*, *in Seoul*)

A negative example comes from the rather debatable situation below: *"Lorillard Inc., the unit of New York-based Loews Corp. that makes Kent cigarettes, stopped using crocidolite in its Micronite cigarette filters in 1956."* PropBank gives these relations: AGT(*Lorillard Inc.*, *use*) and AT-L(*use*, *in Micronite cigarette filters*) However, the conclusion below does not make sense.

AGT(*Lorillard Inc.*, *use*) ∘ AT-L(*use*, *in Micronite cigarette filters*) $\nrightarrow$ AT-L(*Lorillard Inc.*, *in Micronite cigarette filters*)

The reason for this axiom to result in an odd conclusion is because the *"Micronite cigarette filters"* was not really the location for *"use"*; it is indeed the location for *"crocidolite"*.

Similar experiment was run for WordNet and XWN-KB. There are 503 cases where AT-L and all typical syntactic subjects relations share a common concept for semantic composition. We randomly selected one hundred of them and they were all correct, giving hundred per cent precision.

Despite these minor exceptions, the axiom works for the majority of the cases in the experiment, scoring on average high percentage of precision.

6.4 Typical direct object *cluster*

In typical direct object cluster, THM relation can compose with AT-L relation (see Section 5.5). When used PropBank sentences, there are 138 sentences where AT-L and THM relations share a concept. We have randomly selected one hundred, finding ninety-seven positive cases where the axiom produced a valid conclusion, scoring (97/100) = 97 per cent precision. In an example with the sentence *"An exhibition of American design and architecture opened in September in Moscow and will travel to eight other Soviet cities"*, we have THM(*exhibit*, *open*) and AT-L(*open*, *in Moscow*). Then, we have the below axiom that produces a valid instance of a relation.

THM(*exhibit*, *open*) ∘ AT-L(*open*, *in Moscow*) → AT-L(*exhibit*, *in Moscow*)

We used XWN-KB for the same axioms and found 832 cases in which the two relations share a common concept. A similar random check of one hundred revealed only four negative cases, resulting in ninety-six per cent precision.

6.5 is-a *relation*

For case 1 (ISA o AT-L) of AT-L and ISA composition, we have the following example from XWN-KB for the concept *"Tellurian"* which has a gloss of *"an inhabitant of the earth"*. Among the relations that are encoded by the semantic parser are AT-L(*inhabitant*, *earth*) and ISA(*Tellurian*, *inhabitant*). These two relations can be semantically composed by the axiom of case 1 and the conclusion below can be obtained:

ISA(*Tellurian*, *inhabitant*) ∘ AT-L(*inhabitant*, *earth*) → AT-L(*Tellurian*, *earth*)

Similar conclusion can be obtained for case 2 (ISA o AT-L^{-1}). In an example from XWN-KB the concept *"alehouse"* has the gloss *"a tavern where ale is sold"* having the following relations produced by the semantic parser:
ISA(*tavern*, *alehouse*) and AT-L(*sell*, *tavern*). Using the composition axiom, we arrive at the valid instance below, which did not exist before the application of the axiom.

$$\text{AT-L}(sell, tavern) \circ \text{ISA}^{-1}(tavern, alehouse) \rightarrow \text{AT-L}(sell, alehouse)$$

The experiment scored a high precision, as shown in the table. The valid axioms for all four cases are applied to ISA (Hypernymy, Hyponymy) and AT-L relations in WordNet and XWN-KB. We have obtained 1997 conclusive instances with a common shared concept. We randomly checked one hundred and they all made sense, as in Table 13.

6.6 **part-whole** *relation*

In this experiment, for the composition of PW and AT-L, we use annotated AT-L and PW relations in WordNet and XWN-KB.

In the experiment, we encountered 179 cases where the two relations share a common concept. Some of the actual examples are given in Section 5.7.

Of 179, one hundred were checked and all made sense, showing that, using the semantic composition with PW, the precision is hundred per cent.

6.7 **make** *relation*

Among all four cases, only case 4 (AT-L^{-1} o MAK) has the valid conclusion for the composition of the MAK and AT-L relations. We did the experiment in XWN-KB and looked for the composition of case 4. Following are some actual examples. The *"Ketone body"* has the gloss *"a ketone that is an intermediate product of the breakdown of fats in the body; any of three compounds (acetoacetic acid, acetone, and/or beta-hydroxybutyric acid) found in excess in blood and urine of persons with metabolic disorders"*. The gloss has the following relations: AT-L(*breakdown of fats*, *body*) and MAK(*ketone*, *breakdown of fats*). These two relations can compose as in case 4 and produce a valid relation below:

$$\text{AT-L}^{-1}(body, breakdown\ of\ fats) \circ \text{MAK}(breakdown\ of\ fats, ketone) \rightarrow \text{AT-L}^{-1}(body, ketone)$$

In this example, we conclude that there is ketone in the body. The problem with this composition is that although the examples are good and make sense; there are not many examples of this kind in XWN-KB; there are only eight cases overall, all of which make sense.

6.8 Synonymy relation

For the experiment, we used XWN-KB. In one example, the *"Cecil Rhodes"* has a gloss *"British colonial financier and statesman in South Africa; made a fortune in gold and diamond mining; helped colonize the territory now known as Zimbabwe; he* *endowed annual fellowships for Commonwealth and United States students to study at Oxford University (1853-1902)"*. From the gloss we have AT-L(*colonize*, *territory*), and SYN(*territory*, *Zimbabwe*) relations, which can be composed as below axiom:

AT-L(*colonize*, *territory*) ∘ SYN(*territory*, *Zimbabwe*) → AT-L(*colonize*, *Zimbabwe*)

With the axiom, now we have a relation that states the colonization is in Zimbabwe. A different example from XWN-KB would be for the *"spatterdock"*. In its gloss of *"common water lily of eastern and central North America, having broad leaves and globe-shaped yellow flowers; in sluggish fresh or slightly brackish water"*, there is a AT-L(*spatterdock*, *brackish water*) relation stating a location for the *"spatterdock"*, which also has very few synonyms in its synset encoded with the relations SYN(*spatterdock*, *water lily*) and SYN(*spatterdock*, *Nuphar advena*). Thus, the composition can result in more relations with the axiom below:

SYN(*water lily*, *spatterdock*) ∘ AT-L(*spatterdock*, *brackish water*) → AT-L(*water lily*, *brackish water*)

During the thorough experiment in XWN-KB, we encountered 264 cases of AT-L and SYN relations. In a random check of one hundred, only two did not make sense, giving ninety-eight per cent precision.

7 Applications and conclusions

There are many use cases and applications that would benefit from CSR. It may facilitate inferring and common sense reasoning. TE is the task of determining if the meaning of one text is entailed, i.e., can be inferred, by another. The interest for TE has grown in recent years, mainly due to the fact that several NLP applications, such as Question Answering, Information Extraction, and Summarizing can take advantage of semantic inferences to improve their performance. The recognising textual entailment (RTE) challenge focuses on the evaluation of systems for recognizing TE. Consider the task of determining if text **T** entails hypothesis **H**; pair id=38 in RTE3 (Giampiccolo *et al.* 2007)

T: *"Cote d'Ivoire, once a haven of stability in West Africa and the world's top cocoa producer, has been split in two since a failed coup against Gbagbo in September 2002, pitting rebels from the Muslim-dominated north against the Christian-populated south."*

H: *"Cote d'Ivoire is located in Africa."*

Given **T**, the semantic relation AT-L(*Cote d'Ivoire*, *West Africa*) can be obtained. From WordNet, we also have PW_P(*West Africa*, *Africa*). Using case 3 of AT-L and PW composition (see Table 9), we conclude AT-L(*Cote d'Ivoire*, *Africa*). From H, the same semantic relation AT-L(Cote d'Ivoire, Africa) is extracted. We can, therefore, easily conclude that **T** entails **H** by simply matching the semantic relations from **T** and **H**.

Extending this research, other pairs of relations can also be studied under CSR, potentially enriching the overall semantics of text. For example, CAUSE and PURPOSE yield interesting conclusions that enrich the existing semantics of text. Another interesting improvement would be to introduce constraints for axioms that yield plausible relations. For example, one can tag a size restriction on case 4 (a) in AT-L and PW_p composition (Section 5.7), which states that the whole should not be larger than the location of its part for the axiom to be valid transforming this axiom from plausible to valid under a condition. This would be a good solution for the problem we mentioned in AT-L(*Lubbock*, *Texas*) example earlier in this paper. The method can also be used to instantiate axioms that can be used in many other applications; for example, commonsense reasoning.

In summary, we believe that CSR offers great potential for many semantic rich processing applications since it adds more semantics that is otherwise not expressed.

References

Badulescu, A., and Srikanth, M. 2007 (June). Lcc-srn: Lcc's srn system for semeval 2007 task 4. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pp. 215–8, Prague, Czech Republic: Association for Computational Linguistics.

Baker, C. F., Fillmore, C. J., and Lowe, J. B. 1998. The Berkeley FrameNet project. In *Proceedings of ACL/COLING*, Montreal, Canada, pp. 86–90.

Beamer, B., and Girju, R. 2009. Using a bigram event model to predict causal potential. In *10th International Conference on Intelligent Text Processing and Computational Linguistics (CICLING)*, Mexico City, Mexico, pp. 430–41.

Bethard, S., and Martin, J. H. 2008. Learning semantic links from a corpus of parallel temporal and causal relations. In *Proceedings of ACL-08: HLT, Short Papers*, pp. 177–80. Columbus, OH: Association for Computational Linguistics.

Blaheta, D., and Charniak, E. 2000. Assigning function tags to parsed text. In *Proceedings of the 1st North American chapter of the Association for Computational Linguistics conference*, pp. 234–40. San Francisco, CA: Morgan Kaufmann Publishers Inc.

Carreras, X., and Màrquez, L. 2004. Introduction to the conll-2004 shared task: semantic role labeling. In H. T. Ng, and E. Riloff (eds.), *HLT-NAACL 2004 Workshop: Eighth Conference on Computational Natural Language Learning (CoNLL-2004)*, pp. 89–97. Boston, MA: Association for Computational Linguistics.

Chang, D. S., and Choi, K. S. 2006. Incremental cue phrase learning and bootstrapping method for causality extraction using cue phrase and word pair probabilities. *Information Processing & Management* **42**(3): 662–78.

Chen, J., and Rambow, O. 2003. Use of deep linguistic features for the recognition and labeling of semantic arguments. In *Proceedings of the 2003 Conference on Empirical methods in Natural Language Processing*, pp. 41–8. Morristown, NJ: Association for Computational Linguistics.

Culotta, A., McCallum, A., and Betz, J. 2006. Integrating probabilistic extraction models and data mining to discover relations and patterns in text. In *In Proceedings of HLT-NAACL'06*, New York City, pp. 296–303.

Dang, H. T., and Palmer, M. 2005. The role of semantic roles in disambiguating verb senses. In *Proceedings of ACL*, Ann Arbor, Michigan.

Fillmore, C. 1971. Some problems for case grammar. *Monograph Series on Languages and Linguistics* **24**: 35–6.

Giampiccolo, D., Magnini, B., Dagan, I., and Dolan, B. 2007. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL Workshop* *on Textual Entailment and Paraphrasing*, pp. 1–9. Prague, Czech Republic: Association for Computational Linguistics.

Girju, R., Nakov, P., Nastase, V., Szpakowicz, S., Turney, P., and Yuret, D. 2007. Semeval-2007 task 04: classification of semantic relations between nominals. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pp. 13–8. Prague, Czech Republic: Association for Computational Linguistics.

Giuglea, A. M., and Moschitti, A. 2006. Semantic role labeling via FrameNet, VerbNet and PropBank. In *ACL-44: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pp. 929–36. Morristown, NJ: Association for Computational Linguistics.

Golledge, R. G., Gale, N., Pellegrino, J. W., and Doherty, S. 1992. Spatial knowledge acquisition by children: route learning and relational distances. *Annals of the Association of American Geographers* **82**(2): 223–44.

Halsbury, Lord. 1962. Professor Waddington's naturalistic ethics. *Philosophy* **37**(139): 63–6.

Harabagiu, S. 1998. Deriving metonymic coercions from WordNet. In *Proceedings of the Workshop on Usage of WordNet in Natural Language Processing Systems, COLING-ACL'98*, Montréal/Canada, pp. 142–8.

Harabagiu, S., and Moldovan, D. 1998. *WordNet: An Electronic Lexical Database and Some of its Applications. Knowledge processing on extended WordNet*, ch. 17, pp. 684–714. The MIT Press.

Heft, H., and Wohwill, J., 1987. *Handbook of Environmental Pyschology*, pp. 175–204. John Wiley and Sons.

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

Hendrickx, I., Kim, S. N., Kozareva, Z., Nakov, P., Séaghdha, D., Padó, S., Pennacchiotti, M., Romano, L., and Szpakowicz, S. 2010. SemEval-2010 task 8: Multi-Way classification of semantic relations between pairs of nominals. In *Proceedings of the 5th SIGLEX Workshop on Semantic Evaluation (ACL 2010)*, Uppsala, Sweden.

Jackendoff, R. 1983. *Semantics and Cognition*. Cambridge, MA: MIT Press.

Jackendoff, R. 1990. *Semantic Structures*. Cambridge, MA: MIT Press.

Kenny, A. 1966. Intention and purpose. *The Journal of Philosophy* **63**(20): 642–51.

Kingsbury, P., Palmer, M., and Marcus, M. 2002. Adding semantic annotation to the Penn TreeBank. In *In Proceedings of the Human Language Technology Conference*, San Diego, CA.

Levinson, S. C. 1996. Language and space. *Annual Review of Anthropology* **25**: 353–82.

Marcus, M., Santorini, B., and Marcinkiewicz, M. A. 1994. Building a large annotated corpus of English: the penn treebank. *Computational Linguistics* **19**(2): 313–30.

Màrquez, L., Carreras, X., Litkowski, K. C., and Stevenson, S. 2008. Semantic role labeling: an introduction to the special issue. *Computational Linguistics* **34**(2): 145–59.

Mihalcea, R., and Moldovan, D. 2001. eXtended WordNet: progress report. In *Proceedings of NAACL Workshop on WordNet and Other Lexical Resources*, Pittsburgh, PA.

Mihalcea, R., and Shi, L. 2005. Putting pieces together: combining FrameNet, VerbNet and WordNet for robust semantic parsing. In *Proceedings of the Sixth International Conference on Intelligent Text Processing and Computational Linguistics*, Mexico City, Mexico.

Miller, G. A. 1995. WordNet: a lexical database for english. *Communications of the ACM* **38**(11): 39–41.

Miller, G. A., and Johnson-Laird, P. N. 1976. *Language and Perception*. Cambridge, MA: Belknap Press of Harvard University Press.

Moldovan, D., Badulescu, A., Tatu, M., Antohe, D., and Girju, R. 2004. Models for the semantic classification of noun phrases. In *In HLT-NAACL 2004: Workshop on Computational Lexical Semantics*, Boston, MA, pp. 60–7.

Palmer, M., Gildea, D., and Kingsbury, P. 2005. The proposition bank: an annotated corpus of semantic roles. *Computational Linguistics* **31**(1): 71–106.

Piaget, J. 1952. *The Child's Conception of Number*. London: Routledge & Kegan Paul. Translated by C. Gattegno and F.M. Hodgson.

Ponzetto, S. P., and Strube, M. 2006. Semantic role labeling for coreference resolution. In *EACL '06: Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics: Posters & Demonstrations*, pp. 143–6. Morristown, NJ: Association for Computational Linguistics.

Pradhan, S. S., Ward, W. H., Hacioglu, K., Martin, J. H., and Jurafsky, D. 2004. Shallow semantic parsing using support vector machines. In D. M. S. Dumais, and S. Roukos (eds.) *HLT-NAACL 2004: Main Proceedings*, pp. 233–40, Boston, MA: Association for Computational Linguistics.

Srikumar, V., Reichart, R., Sammons, M., Rappoport, A., and Roth, D. 2008. Extraction of entailed semantic relations through syntax-based comma resolution. In *Proceedings of ACL-08: HLT*, pages 1030–8. Columbus, OH: Association for Computational Linguistics.

Szpakowicz, S., and Barker, K. 1995. Interactive semantic analysis of clause-level relationships. In *Proceedings of the Second Conference of the Pacific Association for Computational Linguistics*, Brisbane, pp. 22–30.

Tatu, M. 2007. *Intentions in Text and Semantic Calculus*. PhD. thesis, University of Texas at Dallas, Richardson, TX.

Tatu, M., and Moldovan, D. 2007. Cogex at rte 3. In *In Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, Prague, Czech Republic, pp. 22–7.

Turney, P. D. 2006. Expressing implicit semantic relations without supervision. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pp. 313–20. Sydney, Australia: Association for Computational Linguistics.

Whiteman, M. 1967. *Philosophy of Space and Time and the Inner Constitution of Nature; a Phenomenological Study*. London: Allen & Unwin.

Winston, M. E., Chaffin, R., and Herrmann, D. 1987. A taxonomy of part-whole relations. *Cognitive Science* **11**(4): 417–44.

Wu, D. and Fung, P. 2009. Semantic roles for SMT: a hybrid two-pass model. In *NAACL '09: Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*, pp. 13–6. Morristown, NJ: Association for Computational Linguistics.