Finding Associations between People

Eduardo Blanco and **Dan Moldovan**
Lymba Corporation
Richardson, TX 75080 USA
{*eduardo,moldovan*}@*lymba.com*

Abstract

Associations between people and other concepts are common in text and range from distant to close connections. This paper discusses and justifies the need to consider subtypes of the generic relation ASSOCIATION. Semantic primitives are used as a concise and formal way of specifying the key semantic differences between subtypes. A taxonomy of association relations is proposed, and a method based on composing previously extracted relations is used to extract subtypes. Experimental results show high precision and moderate recall.

1 Introduction

People interact with each other and in doing so foster associations. These associations range from very weak to strong connections. For example, imagine John takes a flight from Rome to New York City and his wife is waiting for him at his destination. Among others, the above scenario indicates that John is associated with any other person in the same flight (passengers and crew) and with his wife. The former association is transient, whereas the latter is lasting.

People are associated with other people (relatives, participants in conversations, employees of the same company, etc.), organizations (schools they attend, clubs they are members of, etc.), locations (places where they were born, live, own property, etc.) and others entities. Associations might be grounded on a spatial context (e.g., two people are associated if they were born in the same country), time overlap (e.g., two people are associated if they belong to the same generation), similar attributes (e.g. similar education or affiliations) or other criteria. Where to draw the line between useful associations and associations too remote to be taken into account is an open question. We believe that the application for which associations are being extracted for defines the *relevant* associations and their granularity.

In this paper, we investigate associations between people. We propose a set of semantic primitives to characterize subtypes of association and define a taxonomy of ASSOCIATION relations. Primitives effectively capture elemental semantic properties of each subtype of ASSOCIATION, making clear the differences in meaning between them and justifying the claim that considering the generic relation ASSOCIATION is insufficient. A two-step approach for extracting subtypes of association grounded on composing previously extracted relations is proposed and experimental results discussed.

2 People Associations

People are associated with each other in a variety of ways, ranging from fairly weak (and yet useful for certain tasks) relations such as SHARED-ORIGIN and stronger links such as VERBAL-COMMUNICATION. Any two individuals born in the same country share their origin even though they might have never met. However, if two people communicate, they interact with each other and (presumably) interchange some information. Simply using ASSOCIATION for any kind of connection between two individuals is too coarse for most applications; more specific relations are needed.

The work on ASSOCIATION presented in this section somewhat resembles the work done by Winston, Chaffin, and Herrmann (1987) on PART-WHOLE. We aim at differentiating and explaining subtypes of ASSOCIATION based on semantic grounds. This allows us to make formal and clear distinctions between subtypes, as well as justify why considering only ASSOCIATION is too broad. Simply put, ASSOCIATION is too generic to be dealt with. For example, SHARED-ORIGIN is intuitively transitive (if x has the same origin than y and y has the same origin than z, x and z share the same origin), but VERBAL-COMMUNICATION is not (if x communicates verbally with y and y communicates verbally with z, x does not necessarily communicate verbally with z). Some associations can be directly linked to a joint event between the associates (e.g., VERBAL-COMMUNICATION), others are due to two different events sharing a common entity (e.g., SHARED-ORIGIN). Whereas there is a strong link between two people that communicate verbally, two people born in the same location (e.g., city, country) are associated in a (potentially) fortuitous manner.

2.1 Semantic Primitives for Subtypes of Association

We define subtypes of ASSOCIATION using semantic primitives. Primitives specify elemental semantic properties between the arguments of relations and help define them. For example, the primitive *common location* indicates whether a relation requires the arguments to be at a shared location.

SPOKEN-COMMUNICATION holds this primitive (two people must be in the same location in order to talk to each other) and PHONE-COMMUNICATION does not (two people in different locations can communicate by phone).

Each relation takes a value for each primitive from the set $V = \{-, 0, +\}$, where '$-$' indicates that the primitive does not hold, '$+$' that it holds and '0' that it does not apply. In this paper, we use semantic primitives exclusively to define relations. However, we have shown that an algebra for composing primitives can be used to automatically obtain inference axioms for composing any set of semantic relations (Blanco and Moldovan 2011).

We propose a set of nine primitives. These primitives define some semantic dimensions around which any association can be defined. We purposely avoid primitives used elsewhere but not helpful to define the subtypes of ASSOCIATION we work with (e.g., *near* (Huhns and Stephens 1989)). First, we introduce primitives defined by us to differentiate subtypes of ASSOCIATION ordered by decreasing generality and discrimination power (primitives 1–6). Second, we present primitives previously proposed to define other relations (primitives 7–8 (Winston, Chaffin, and Herrmann 1987), 9 (Blanco and Moldovan 2011)). In the definitions below, R denotes a subtype of ASSOCIATION, *x* the first argument of R and *y* the second argument:

1. **Active/Passive.** If R is due to an action performed by *x*, R is active ($+$); otherwise passive ($-$). For example, SHARED-ORIGIN is passive (i.e., *x* and *y* can share their origins without *x* performing any action for that association to hold) and any type of COMMUNICATION is active (i.e., in order for *x* to communicate with *y*, *x* must perform some communicating action).
2. **Direct/Indirect.** If R requires an explicit activity between *x* and *y*, direct ($+$); otherwise indirect ($-$). For example, any COMMUNICATION and ECONOMIC-TRADE is direct (the communicating action or economic trade occurs between *x* and *y*) and SHARED-ORIGIN is indirect (*x* and *y* were (likely) born in separate events).
3. **Common Event.** *x* and *y* are involved in a common event for R to hold. For example, any COMMUNICATION involves the two participants in a single event (both *x* and *y* participate in the act of communicating) and SHARED-ORIGIN does not.
4. **Common Location.** *x* and *y* must be at a common physical location for R to hold. For example, SPOKEN-COMMUNICATION holds this primitive and ELECTRONIC-COMMUNICATION does not.
5. **Material/Immaterial.** If a material object is transferred from *x* to *y* in order for R to hold, material ('$+$', e.g., SENDS-TO). If an immaterial object is transferred (ideas, thoughts, etc.), immaterial ('$-$', e.g., TEACHES). If no object is transferred, this primitive does not apply ('0', e.g., SHARED-ORIGIN).
6. **Hierarchical.** *x* is superior in a power hierarchy to *y*. For example, teachers are higher than students (TEACHES), but COMMUNICATES-WITH and MEETS do not require a hierarchy difference between *x* and *y*.
7. **Functional.** *x* is in a specific spatial or temporal position with respect to *y* for the association to hold. For example, LETTER-COMMUNICATION does not hold this primitive (*x* and *y* do not need to be in any specific location / time to communicate by letter) and PHONE-COMMUNICATION does (*x* and *y* must be by a phone *during* the call).
8. **Homeomerous.** *x* must be the same kind of thing as *y*. For example, only two people can COMMUNICATE or be associated by an ECONOMIC-TRADE (same kind, '$+$'), but a person and the school he attends are not (semantically) of the same kind (ATTENDS-SCHOOL, '$-$').
9. **Universal.** *x* is associated via R to *y* at all times; the association is universally true. For example, two people either have or do not have SHARED-ORIGIN, but they might be AFFILIATED only during a certain period of time.

These primitives are not guaranteed to be the best for any given set of subtypes of ASSOCIATION. They are not independent of each other either. For example, if any relation holds *common location*, it must also hold *functional*. In other words, if two people must be in a common location for R to hold, they must also be in a specific spatial or temporal position. Note that the opposite is not true: if two people are associated through PHONE-COMMUNICATION, they must be in a specific time frame (while the communication takes place) by a phone (*functional* holds), but they are likely to be in different locations (*common location* does not hold).

The specialized primitives for subtypes of association (primitives 1–6) are fairly general. This is shown in Table 1, where several subtypes typically share the same value ('$-$', '0' or '$+$') for the same primitive. Primitives 7–9 are less general when applied to subtypes of association (e.g., only SHARED-ORIGIN holds universal), but they have been justified and proven useful when dealing with other semantic relations (Winston, Chaffin, and Herrmann 1987; Huhns and Stephens 1989).

2.2 A Taxonomy of Associations

Table 1 presents the taxonomy of association relations we propose. There are seven main types of association: COMMUNICATES-WITH, MEETS, JOINT-WORK, ECONOMIC-TRADE, INSTRUCTS, AFFILIATIONS and SPATIALLY-ASSOCIATED. COMMUNICATES-WITH is further divided into two more specific subtypes: WRITTEN-COMMUNICATION and VERBAL-COMMUNICATION.

Note that the taxonomy is supported by the values each subtype takes for the primitives. Associations under SPATIALLY-ASSOCIATED are the only ones that are passive; associations under AFFILIATIONS are the only ones active and indirect; subtypes of VERBAL-COMMUNICATION are the only ones that are immaterial, hold functional and do not hold hierarchical; subtypes of WRITTEN-COMMUNICATION are the only ones that do not hold common location or functional, and are either material or immaterial, and so on.

3 Approach

We propose a two-step approach to extract subtypes of ASSOCIATION. First, we use an existing semantic parser to extract basic semantic relations from text. This parser is trained using standard supervised techniques (Section 4.1) and any

ASSOCIATION	1: Active/Passive	2: Direct/Indirect	3: Common Event	4: Common Location	5: Material/Immaterial	6: Hierarchical	7: Functional	8: Homeomerous	9: Universal
COMMUNICATES-WITH									
WRITTEN-COMMUNICATION									
LETTER-COMMUNICATION	+	+	+	−	+	−	−	+	−
ELECTRONIC-COMMUNICATION	+	+	+	−	−	−	−	+	−
VERBAL-COMMUNICATION									
PHONE-COMMUNICATION	+	+	+	−	−	−	+	+	−
SPOKEN-COMMUNICATION	+	+	+	+	−	−	+	+	−
MEETS	+	+	+	+	0	−	+	+	−
JOINT-WORK									
EMPLOYS	+	+	−	−	0	+	−	+	−
WORKS-WITH	+	+	−	−	0	−	−	+	−
SHARES-TASK-WITH	+	+	+	−	0	−	+	+	−
ECONOMIC-TRADE									
SENDS-TO	+	+	+	−	+	−	+	+	−
SELLS-TO	+	+	+	−	+	−	+	+	−
TRANSFERS-TO	+	+	+	−	+	−	+	+	−
INSTRUCTS									
TEACHES	+	+	+	+	−	+	+	+	−
ATTENDS-SCHOOL	+	+	−	+	0	−	+	−	−
TEACHES-SCHOOL	+	+	−	+	0	−	+	−	−
AFFILIATIONS									
AFFILIATED	+	−	−	−	0	−	−	+	−
MEMBER-OF	+	−	−	−	0	−	−	−	−
SPATIALLY-ASSOCIATED									
SHARED-ORIGIN	−	−	−	−	0	−	−	+	+
SHARED-LOCATION	−	−	−	+	0	−	+	+	−
SHARED-DESTINATION	−	−	−	−	0	−	+	+	−

Table 1: Taxonomy of ASSOCIATION relations and values for semantic primitives. Primitives 1–6 are introduced by us to differentiate subtypes of ASSOCIATION; primitives *functional* and *homeomerous* were proposed by Winston, Chaffin, and Herrmann (1987) to differentiate subtypes of PART-WHOLE.

tool designed to extract semantic relations could be used instead. Second, we instantiate inference axioms that take as premises chains of basic relations extracted by the parser and output a new relation between the ends of the chain.

We denote $R(x, y)$ a semantic relation R holding between x and y. For example, AGENT(*John*, *sent*) is interpreted *John is the agent of sent*. The composition operator '$\circ$' is used to chain relations. An inference axiom consists of a chain of relations (premises) and a conclusion; the symbol '$\rightarrow$' is used to separate premises and conclusion. For example, $R_1(x, y) \circ R_2(y, z) \rightarrow R_3(x, z)$ indicates that given a chain of relations formed by $R_1(x, y)$ and $R_2(y, z)$, one can infer $R_3(x, z)$. We write this axiom $R_1 \circ R_2 \rightarrow R_3$ to save space. In order to chain two relations, they must have an argument in common. We denote R^{-1} the inverse relation of R, which corresponds to relation R with swapped arguments (if we have AGENT(*John*, *sent*), we also have AGENT^{-1}(*sent*, *John*)).

For example, given *John sent a letter to Bill*, our semantic parser extracts AGENT(*John*, *sent*), THEME(*letter*, *sent*) and RECIPIENT(*Bill*, *letter*). Instantiating axiom $\text{AGENT} \circ \text{THEME}^{-1} \circ \text{RECIPIENT}^{-1} \rightarrow$ WRITTEN-COMMUNICATION, we obtain the new relation WRITTEN-COMMUNICATION(*John*, *Bill*). A paraphrase of the above inference is *the agent of a sending event and the recipient of the object sent communicate in writing*. In this paper, we work with manually extracted axioms.

3.1 Constraining Axioms

Generic axioms like the above are useful in certain circumstances, but it seems obvious that restrictions on the arguments of premises are needed. Otherwise, we would generate a lot of false positives. For example, (1) *John sent flowers to Bill* has exactly the the same chain of AGENT and THEME^{-1} as (2) *John sent a letter to Bill*, but the object sent

Name	Synsets
receiving	receive.v.{1, 3}
meeting	meeting.n.1, gathering.n.{1, 2}, meet.v.{1, 2, 3}, gather.v.2
attending	attend.b.1, be.v.3,
speaking	speak.v.{1, 2, 4}, say.v.{1, 6}, tell.v.{2, 3}, affirm.v.3, admit.v.1, ask.v.{1, 2, 3, 6}
transferring	give.v.{3, 17}, transfer.v.{3, 5}, transport.v.{1, 2, 4}, move.v.2

Table 2: Lexical chains and relevant WordNet synsets.

The terrorists sent letter bombs a few years ago to newspaper offices in New York City and Washington, D.C., [. . .]

AGENT(*The terrorists*, *sent*) THEME(*letter bombs*, *sent*) IS-A(*letter bomb*, *letter*) IS-A(*letter bomb*, *bomb*) PURPOSE(*bombs*, *letter bomb*) TEMPORAL(*a few years ago*, *sent*) LOCATION(*to newspaper [. . .], D.C.*, *sent*) MAKE(*newspaper*, *offices*) RECIPIENT(*newspaper [. . .], D.C.*, *letter bombs*) IS-A(*newspaper offices*, *offices*) LOCATION(*New York City*, *newspaper offices*) LOCATION(*Washington, D.C.*, *newspaper offices*) PART-WHOLE(*Washington*, *D.C*) LOCATION(*D.C.*, *Washington*)

Figure 1: Example of semantic relations our semantic parser aims at extracting.

(i.e., *flowers* versus *letter*) should block inferring WRITTEN-COMMUNICATION using the above axiom in (1) (but not in (2)). We consider two kinds of restrictions: named entities (ne) and lexical chains (lc). We use an in-house named entity recognizer and word sense disambiguator.

Lexical Chains are sequences of semantically related words that link two concepts. We define lexical chains along the semantic relations in WordNet (Fellbaum 1998). The strength of the semantic connection between two concepts depends on the number of lexical chains (the more chains the stronger), lengths (the shorter the stronger) and relations involved (HYPONYM is stronger than PART-WHOLE). We use the weighting scheme proposed by Novischi and Moldovan (2006).

Lexical chains are used to define constraints for axioms. They allow us to find different lexical realizations of the same semantic argument. For example, concepts reachable from *meeting.n.1* via lexical chains include *conclave.n.1*, *reunite.v.1* and *gather.v.2*. Table 2 presents some of the lexical chains we consider in the axioms presented in this paper. Following the table, we consider a relation argument a *meeting* if the sum of the weights of the lexical chains between the argument and concepts *meeting.n.1*, *gathering.n.1*, *gathering.n.2*, *meet.v.1*, *meet.v.2*, *meet.v.3* or *gather.v.2* is above a threshold.

4 Implementation

4.1 Semantic Parser

We use an in-house semantic parser that given plain text extracts binary semantic relations (AGENT, THEME, etc.). Our parser extracts not only semantic roles (relations between a verb and its arguments), but also relations between nominals (*door knob*: PART-WHOLE(*door knob*, *door*)) or adjectival phrases (*cat in the tree*: LOCATION(*tree*, *cat*)). Details about the parser can be found in (Moldovan and Blanco 2012), we only briefly describe it here. An example of text and the relations the parser aims at extracting is provided in Figure 1.

The parser is trained using a mixture of annotations from PropBank, FrameNet, NomBank, SemEval competitions and annotations done by us. It is implemented following eight main steps: pre-processing (POS tagging, NER, syntactic parsing, WSD, co-reference resolution), bracketer (simplification of parse trees), argument identification (concept pairs likely to encode a relation are identified based on syntactic patterns), domain and range filtering (plausible relations are proposed based on semantic classes), grouping (pairs are clustered into eight generic syntactic patterns), feature extraction, classifiers and conflict resolution (final relations are decided upon ranking the classifiers output).

4.2 Instantiating Axioms

Instantiating axioms is as simple as finding a chain of relations matching the premises and infer the conclusion if the restrictions are met. Let us consider an axiom $R_1(x, y) \circ R_2(y, z) \rightarrow R_3(x, z)$ with three restrictions: ne(x) = *human*, ne(z) = *human* and lc(y) = *receiving*. We loop through all the relations extracted by the semantic parser and for each instantiation of the premises, enforce the restrictions. Let us denote an instantiation $R_1(a, b) \circ R_2(b, c)$, where *a*, *b* and *c* are actual concepts. If both *a* and *c* are *human* according to the named entity recognizer, the first two restrictions are fulfilled. The next step is to check if *b* is a *receiving* concept. We obtain the weights of the lexical chains between *b* and all the synsets related to *receiving* according to Table 2. This is done by navigating along WordNet relations until a predefined minimum weight is reached (the more relations involved, the lesser the weight). If the total weight (the sum of the weights of all lexical chains) is above the threshold, all restrictions are fulfilled and $R_3(a, c)$ inferred. Otherwise, the axioms is not instantiated.

5 Experiments and Results

The experiments reported in this section were done using a corpus of reports within the law enforcement domain. This corpus was divided into two splits: development and test. The development split was used to define axioms. The test split was manually annotated with subtypes of association. Precision and recall are calculated comparing gold annotations with relations obtained instantiating inference axioms.

Axiom 1	SOURCE(*w*, *x*) $\circ$ THEME(*x*, *y*) $\circ$ AGENT^{-1}(*y*, *z*) $\rightarrow$ WRITTEN-COMMUNICATION(*w*, *z*)
Restrictions	`ne`(w) = *human*, `ne`(z) = *human*, `lc`(x) = *letter*, `lc`(y) = *receive*
Interpretation	The source of a letter and the agent of receiving the letter communicate in writing.
Example	During that period, two prior [letters]$_x$ from [John]$_w$ were [received]$_y$ by [Bill]$_z$.
	— SOURCE(*John*, *letters*) $\circ$ THEME(*letters*, *received*) $\circ$ AGENT^{-1}(*received*, *Bill*) $\rightarrow$ WRITTEN-COMMUNICATION(*John*, *Bill*)
Axiom 2	AGENT(*x*, *y*) $\circ$ THEME^{-1}(*y*, *z*) $\rightarrow$ MEETS(*x*, *z*)
Restrictions	`ne`(x) = *human*, `ne`(z) = *human*, `lc`(y) = *meeting*
Interpretation	Agents meet the themes of meetings.
Example	[John]$_x$ allegedly [met]$_y$ with [Bill]$_z$ last time one year ago.
	— AGENT(*John*, *met*) $\circ$ THEME^{-1}(*met*, *Bill*) $\rightarrow$ MEETS(*John*, *Bill*)
Axiom 3	AGENT(*w*, *x*) $\circ$ THEME^{-1}(*x*, *y*) $\circ$ LOCATION(*y*, *z*) $\rightarrow$ MEETS(*w*, *z*)
Restrictions	`ne`(w) = *human*, `ne`(z) = *human*, `lc`(x) = *attending*, `lc`(y) = *meeting*
Interpretation	Someone attending a meeting which was the location of another person meets that person.
Examples	[John]$_w$ [attended]$_x$ a [meeting]$_y$ at which [Bill]$_z$ was also present.
	— AGENT(*John*, *attended*) $\circ$ THEME^{-1}(*attended*, *meeting*) $\circ$ LOCATION(*meeting*, *Bill*) $\rightarrow$ MEETS(*John*, *Bill*)
Axiom 4	AGENT(*w*, *x*) $\circ$ TOPIC^{-1}(*x*, *y*) $\circ$ THEME^{-1}(*y*, *z*) $\rightarrow$ MEETS(*w*, *z*)
Restrictions	`ne`(w) = *human*, `ne`(z) = *human*, `lc`(x) = *speaking*, `lc`(y) = *meeting*
Interpretation	People meet who they say they meet.
Example	[John]$_w$ [admits]$_x$ having [met]$_y$ [Bill]$_z$ a few times, but claims not to know him well.
	— AGENT(*John*, *admits*) $\circ$ TOPIC^{-1}(*admits*, *met*) $\circ$ THEME^{-1}(*met*, *Bill*) $\rightarrow$ MEETS(*John*, *Bill*)
Axiom 5	AGENT(*w*, *x*) $\circ$ PURPOSE^{-1}(*x*, *y*) $\circ$ THEME^{-1}(*y*, *z*) $\rightarrow$ MEETS(*w*, *z*)
Restrictions	`ne`(w) = *human*, `ne`(z) = *human*, `lc`(y) = *meet*
Interpretation	People meet the people they intend to meet.
Example	[John]$_x$ [arrived]$_y$ in Rome, Italy, to [meet]$_y$ with [Bill]$_z$.
	— AGENT(*John*, *arrived*) PURPOSE^{-1}(*arrived*, *meet*) $\circ$ THEME^{-1}(*meet*, *Bill*) $\rightarrow$ MEETS(*John*, *Bill*)
Axiom 6	AGENT(*w*, *x*) $\circ$ THEME^{-1}(*x*, *y*) $\circ$ RECIPIENT^{-1}(*y*, *z*) $\rightarrow$ TRANSFERS-TO(*w*, *z*)
Restrictions	`ne`(w) = *human*, `ne`(z) = *human*, `lc`(y) = *transferring*
Interpretation	Agents transfer goods to the recipients of the objects they transfer.
Example	[He]$_w$ and [John]$_{w'}$ [transported]$_x$ [$7,000]$_y$ to [Mary]$_z$.
	— AGENT(*He*, *transported*) $\circ$ THEME^{-1}(*transported*, *$7,000*) $\circ$ RECIPIENT^{-1}(*$7,000*, *Mary*) $\rightarrow$ TRANSFERS-TO(*He*, *Mary*)
	— AGENT(*John*, *transported*) $\circ$ THEME^{-1}(*transported*, *$7,000*) $\circ$ RECIPIENT^{-1}(*$7,000*, *Mary*) $\rightarrow$ TRANSFERS-TO(*John*, *Mary*)

Table 3: Sample of inference axioms and examples.

A subset of inference axioms and examples are depicted in Table 3. Each axiom can be seen as a rule indicating that whenever a semantic path between two concepts is found, they are related by a subtype of association. The premises indicate the semantic path and the conclusion the subtype of association. For example, axiom 1 could be interpreted as *if there is a path* SOURCE*(w, x)* $\circ$ THEME*(x, y)* $\circ$ AGENT^{-1}*(y, z) (and constraints are fulfilled), then* WRITTEN-COMMUNICATION*(w, z) holds*.

In Table 3, axiom 1 infers WRITTEN-COMMUNICATION, axioms 2–5 MEETS and axiom 6 TRANSFERS-TO. Note that axiom 6 can be instantiated twice in the statement provided as example and therefore two different relations TRANSFER-TO (TRANSFER-TO(*He*, *Mary*) and TRANSFER-TO(*John*, *Mary*)) are inferred. Some inferences (e.g., TRANSFER-TO(*He*, *Mary*)) need an extra step of co-reference resolution to be useful, but we do not show it.

The total number of axioms defined is 158. Following the steps in Section 4.2, we instantiate them over the output of the semantic parser in the test split. Overall results are shown in Table 4. It is worth noting that the 238 inferred associations (ignored by the semantic parser) add 2.34% semantic relations in relative terms on top of the 10,141 relations extracted by the parser.

Number of sentences	1,000
Number of basic semantic relations parser outputs	10,141
Number of inferences	238
Precision	0.81
Recall	0.60
F-measure	0.69

Table 4: Overall results on the test split. Each axiom instantiation leads to an inference.

Most errors are due to the semantic parser: instantiating an axiom based on incorrect premises almost always implies an incorrect inference. Recall is low because the 158 axioms do not account for all possible ways to express an association. First, axioms are limited by the relations outputted by the semantic parser (and there is no guarantee that all associations can be obtained by linking chains of relations the parser extracts). Second, they were manually obtained after examination of the development set. Thus, even if the chain does exist, there is no guarantee that an axiom capable of inferring it was proposed.

6 Related Work

Semantic primitives have been used before to define semantic relations under the names of *relation elements*, *deep structure*, *aspects* and *primitives*. The number of primitives varies depending on the relation inventory. Chaffin and Herrmann (1987) differentiate a set of 31 relations clustered in five groups (contrast, similars, class inclusion, case-relations, part-whole) with 30 *relation elements* also clustered in five groups (elements of intensional force, dimension elements, elements of agreement, propositional elements, elements of part-whole inclusion). Winston, Chaffin, and Herrmann (1987) use 3 *relation elements* (functional, homeomerous and separable) to distinguish between six subtypes of PART-WHOLE. Cohen and Losielle (1988) use only 2 *aspects* (*hierarchical* and *intangible*) and Huhns and Stephens (1989) extend previous work and consider a set of 10 *semantic primitives*.

Work on extracting semantic relations from text is vast. To the best of our knowledge, there is no specific work on defining and extracting associates of people. Supervised relation extraction has matured in recent years, especially semantic role labeling (Gildea and Jurafsky 2002), i.e., relations between a verb and its arguments. There has been work on CAUSE (Chang and Choi 2006), PART-WHOLE (Girju, Badulescu, and Moldovan 2006) and many others. Turney (2006) discusses similarity of semantic relations (attributional and relational) and presents LRA (Latent Relational Analysis), an unsupervised method to measure relational similarity that learns from a large unlabeled corpus.

Unlike previous proposals to extract semantic relations from text, the inference axioms we propose take as their input previously extracted relations and output a new relation. They do not deal directly with text or syntax; they work on a semantic level and instantiating them creates an extra layer of relations previously ignored.

7 Conclusions

In this paper, we have justified the need to consider subtypes of ASSOCIATION and proposed a taxonomy of ASSOCIATION relations. Simply put, a generic ASSOCIATION is too coarse to be concisely defined, and manipulation or reasoning with it seems at the very least challenging.

Semantic primitives capturing key semantic properties of subtypes of ASSOCIATION have been identified. They effectively define the basic ways in which people are associated by specifying, for example, whether they must share a common location or event for the association to hold. The set of primitives presented is neither complete (i.e., incorporating other subtypes of ASSOCIATION might require more primitives) nor perfect for the taxonomy presented (e.g., SENDS-TO and SELLS-TO have the same values for all primitives). When defining primitives, the goal is to propose those which are discriminative, a primitive that only applies to one relation is arguably useless. For example, we could (but have not) add a primitive indicating if the association requires exchange of money between the associates to differentiate SENDS-TO and SELLS-TO. Instead, we use more generic primitives indicating if the association requires the exchange of a material or immaterial object, or no exchange at all.

A method to extract subtypes of association using inference axioms has been proposed. Axioms infer a subtype of ASSOCIATION by composing relations extracted with an existing semantic parser. This approach contrasts with most methods to extract semantic relations: it does not use as input plain text and syntax is not taken into account. We have presented manually obtained axioms after examination of the output of our semantic parser. Semantic constraints are key to filter out false positives. Preliminary evaluation shows high precision and modest recall.

References

Blanco, E., and Moldovan, D. 2011. Unsupervised Learning of Semantic Relation Composition. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (ACL-HLT 2011)*.

Chaffin, R., and Herrmann, D. J. 1987. *Relation Element Theory: A New Account of the Representation and Processing of Semantic Relations*.

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):662–678.

Cohen, P. R., and Losielle, C. L. 1988. Beyond ISA: Structures for Plausible Inference in Semantic Networks. In *Proceedings of the Seventh National conference on AI*.

Fellbaum, C., ed. 1998. *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. The MIT Press.

Gildea, D., and Jurafsky, D. 2002. Automatic Labeling Of Semantic Roles. *Computational Linguistics* 28:245–288.

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

Huhns, M. N., and Stephens, L. M. 1989. Plausible Inferencing Using Extended Composition. In *Proceedings of the 11th international joint conference on Artificial intelligence*.

Moldovan, D., and Blanco, E. 2012. Polaris: Lymba's Semantic Parser. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*.

Novischi, A., and Moldovan, D. 2006. Question Answering with Lexical Chains Propagating Verb Arguments. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the ACL*, 897–904.

Turney, P. D. 2006. Similarity of Semantic Relations. *Computational Linguistics* 32(3):379–416.

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