Annotating Spatial Containment Relations Between Events

Kirk Roberts, Travis Goodwin, Sanda M. Harabagiu

Human Language Technology Research Institute
University of Texas at Dallas
Richardson TX 75080
{kirk,travis,sanda}@hlt.utdallas.edu

Abstract

A significant amount of spatial information in textual documents is hidden within the relationship between events. While humans have an intuitive understanding of these relationships that allow us to recover an object's or event's location, currently no annotated data exists to allow automatic discovery of spatial containment relations between events. We present our process for building such a corpus of manually annotated spatial relations between events. Events form complex predicate-argument structures that model the participants in the event, their roles, as well as the temporal and spatial grounding. In addition, events are not presented in isolation in text; there are explicit and implicit interactions between events that often participate in event structures. In this paper, we focus on five spatial containment relations that may exist between events: (1) SAME, (2) CONTAINS, (3) OVERLAPS, (4) NEAR, and (5) DIFFERENT. Using the transitive closure across these spatial relations, the implicit location of many events and their participants can be discovered. We discuss our annotation schema for spatial containment relations, placing it within the pre-existing theories of spatial representation. We also discuss our annotation guidelines for maintaining annotation quality as well as our process for augmenting SpatialML with spatial containment relations between events. Additionally, we outline some baseline experiments to evaluate the feasibility of developing supervised systems based on this corpus. These results indicate that although the task is challenging, automated methods are capable of discovering spatial containment relations between events.

Keywords: spatial relations, event relations, spatial reasoning

1. Introduction

Events in text implicitly convey spatial information. Implicit spatial inference occurs when no spatial information is explicitly associated with an event. For instance, in the sentence "*The [bombing] victim [died] instantly*", we understand that the spatial bounds on *died* happened within the spatial bounds on *bombing*. Yet this is not directly stated by the contextual evidence. Two further examples of implicit spatial grounding are illustrated by:

(1) [Rafiq Hariri]$_{\text{PARTICIPANT}}$ [submitted]$_{\text{E1}}$ his resignation during a 10-minute [meeting]$_{\text{E2}}$ with the head of state at the [Baabda presidential palace]$_{\text{LOCATION}}$.

(2) As [Egyptian columns]$_{\text{PARTICIPANT}}$ [retreated]$_{\text{E3}}$, Israel's aircraft [attacked]$_{\text{E4}}$ them, using napalm bombs. The [attacks]$_{\text{E5}}$ destroyed hundreds of vehicles and [caused]$_{\text{E6}}$ heavy casualties in [Sinai]$_{\text{LOCATION}}$.

In Example (1), recognizing *Rafiq Hariri* was located in the *Baabda presidential palace* in the event E1 (*submitted*) requires understanding the spatial relationship between E1 and E2 (*meeting*). Moreover, the temporal connective *during* grounds E1 temporally within the time interval associated with E2. While a temporal relation does not always guarantee a spatial relation, in this case an OCCASION discourse relation exists between E1 and E2, allowing the spatial inference to be made as well. Since the temporal relation indicates event E2 temporally contains event E1, the inference that E2 also spatially contains E1 may be made as well. As E2 is contained within the location *Baabda presidential palace*, E1 is also contained within this location by the transitive property. This allows us to draw the inference that E1's participant *Rafiq Hariri* is located at the palace as well. In Example (2), understanding that the *Egyptian columns* are spatially related to *Sinai* requires understanding the spatial relationships between both coreferential and non-coreferential events. In this example, the event corresponding to E3 (*retreated*) refers to a motion whose source is the same location as the event E4 (*attacked*). Event E4 is coreferential with event E5 (i.e., they are both *mentions* of the same event). Further, E5 has the result event E6 (*caused*), which is located in *Sinai*. These two examples show that spatial containment relations between events can be inferred by relying on many discourse phenomena, including coreference and temporal relations.

Using these types of relations to determine the location of events and their participants fits into the larger work of recovering implicit information (Palmer et al., 1986), where semantically relevant information is found outside an object's syntactic scope. The difficulty with a generalist approach to recovering long-distance semantic arguments, however, is that there is no one set of relations that describes how all implicit information can be recovered. Rather, the method for recovery depends on the type of information being sought. Temporal information, for instance, can be recovered through a very different set of relations than manner or purpose information. This work is part of an effort to create one such method for acquiring implicit spatial information through spatial containment relations between events. Currently, we are unaware of any linguistic resource of manually annotated spatial containment relations between events. We expect the ability to automatically identify spatial containment relations between events will improve the performance of generalized implied semantic role labelers.

In this paper, we describe our method for building such a resource. We consider five basic spatial containment relations between events:

1. SAME: Two events E1 and E2 have indistinguishable spatial bounds.
2. CONTAINS: Either E1's spatial bounds contain E2 or vice versa (this is a directed relation).
3. OVERLAPS: E1 and E2 share partial spatial bounds but neither is a sub-set of the other.
4. NEAR: E1 and E2 do not share spatial bounds but they are within relative proximity of each other.
5. DIFFERENT: E1 and E2 have distinguishably different spatial bounds.

Annotation of all five types of containment relations is performed on SpatialML (Mani et al., 2008).
The remainder of this paper is organized as follows. Section 2 outlines related work in event relations and implicit information recovery. Section 3 discusses our annotation schema, its strengths and weaknesses, as well as our guidelines for annotators. Section 4 describes the process for the creation of our corpus, provides analysis of the annotated documents, and describes the baseline experiments we performed to determine the types of linguistic processing necessary for the automatic recognition of our spatial relations. Finally, Section 5 summarizes our work and proposes future directions for research.

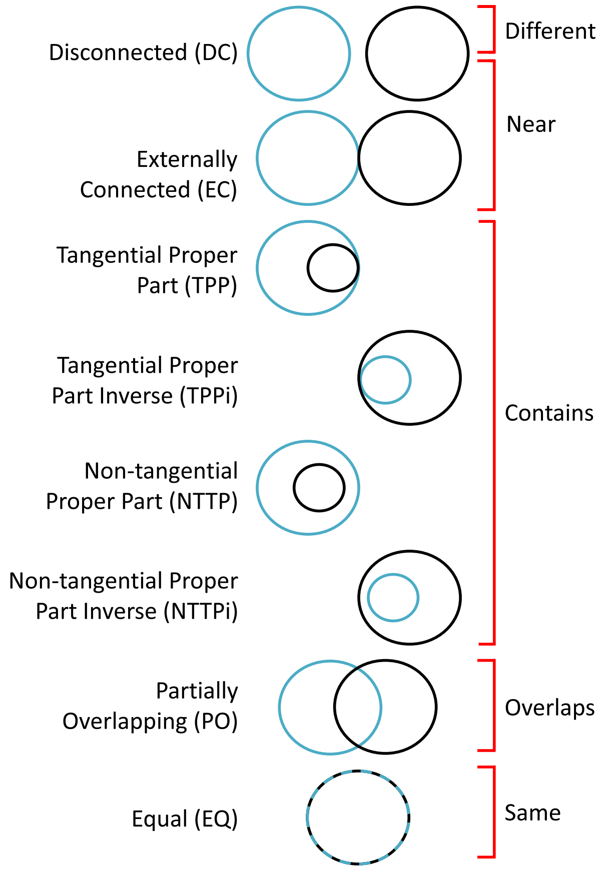

Figure 1: Comparison of our spatial containment relations with RCC-8.

2. Related Work

Event relations represent important knowledge that can be distilled from documents, contributing to discourse and semantic processing, as well as general comprehension of textual information. Among the first to tackle inter-event relation recognition were the researchers that developed TimeML (Pustejovsky et al., 2003a) and its annotated TimeBank corpus (Pustejovsky et al., 2003b). Many event relations follow from discourse theory (Hobbs, 1985b), yielding relations such as causation (Do et al., 2011), coreference (Chen et al., 2009; Bejan and Harabagiu, 2010), and temporal ordering (Chambers and Jurafsky, 2008). Along with these types of inter-event relations, spatial relations between pairs of events allow us to better understand the knowledge that is derived from dependencies between events. In this work, we classify relations according to their inferred spatial relationships, which receive comparably little attention in discourse processing.
Several models of spatial representation in text have been considered, such as ISO-Space (Pustejovsky et al., 2011), SpatialML (Mani et al., 2008), and STML (Pustejovsky and Moszkowicz, 2008). However, the primary goal of these models and their corresponding annotated corpora is to capture spatial relationships explicitly stated in text or the handling of specific sub-classes of events such as motion events. None of these models consider implicit spatial relations between events.
SpatialML in particular is designed to represent spatial locations, largely geographic locations and culturally-relevant landmarks referred to with the PLACE tag. PLACES are disambiguated when necessary (e.g., to differentiate between Paris, France and Paris, Texas). Further, SpatialML includes two types of relations. The PATH relation expresses a spatial trajectory (e.g., "*[northwest] of the capital [New Delhi]*"). The LINK relation expresses containment (e.g., "*a [well] in [West Tikrit]*"). While the PATH relation could be used to describe the spatial relationship between events (e.g., "*the [evacuations] were occurring south of the [riots]*"), this kind of relation is not commonly encountered. The LINK relation, however, is directly analogous to the relations we studied in this research, and we compare the types of containment relations in SpatialML to our own relation types below.

A comparison between our relations and the well-known RCC-8 specification (Randell et al., 1992) relations is shown in Figure 1. Notably, we combine four relations into CONTAINS and add a NEAR relation. SpatialML makes similar simplifications, and is identical to our relations with the exception that it specifies an extended connection relation similar to EC from RCC-8. This, however, is commonly expressed when two locations border each other (e.g., "*the border between [Lebanon] and [Israel]*"). Since the spatial boundary of an event is almost always underspecified, this relation is unlikely to be conveyed in implicit event relations and we thus omit *ec* from consideration.

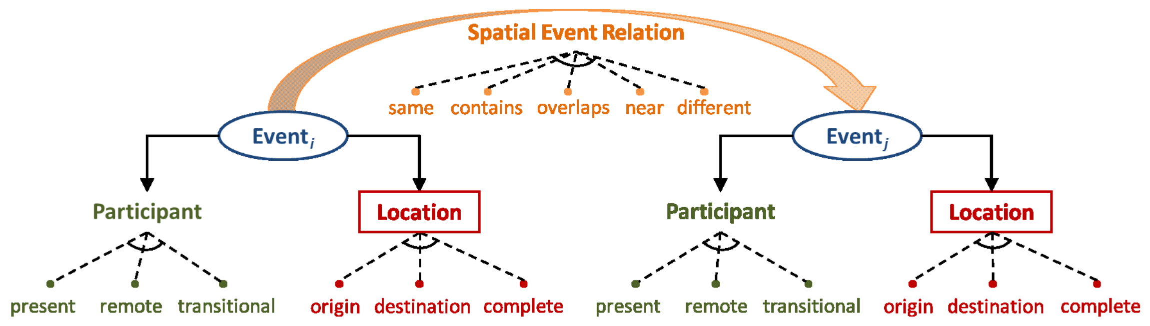

Figure 2: Spatial containment schema. Note that only the CONTAINS relation is directed, all other spatial containment relations are undirected.

3. Annotation Schema and Guidelines

3.1. Definitions

We define an *event* using the TimeML (Pustejovsky et al., 2003a) convention: a situation that happens or occurs. Event *mentions* are words or phrases that denote events in natural language documents and are limited to tensed verbs, event nominals, and stative modifiers. For brevity, we refer to mentions simply as events unless the context requires further clarification.

For LOCATIONS, we follow the SpatialML convention, which includes both named (e.g., Japan) and nominal (e.g., village) toponyms. As with SpatialML, our LOCATIONS could be extended beyond geographic entities to include other types of locations such as biological markers. Our spatial relations would be valid for such domains as well. However, our annotated corpus is limited to the newswire domain, so we focus on geographic locations. Note that not all locations are necessarily event locations (e.g., in "*the United States [fought] in Afghanistan*", *Afghanistan* is a LOCATION but *United States* is a PARTICIPANT).

We limit our definition of PARTICIPANT to persons, organizations, and physical objects to guarantee all participants have spatial properties. To determine which PARTICIPANTS and LOCATIONS to associate with an event, we limit the annotators to the syntactic scope of the event. In other words, the scope expected by a semantic role labeler.

3.2. Schema

Our annotation schema is illustrated in Figure 2. Events are linked to explicit locations in their syntactic scope. There are 3 LOCATION sub-types: COMPLETE, ORIGIN, and DESTINATION. The COMPLETE sub-type indicates that the location identifies the entire spatial bound of the event: it starts and ends at this location without leaving. The ORIGIN and DESTINATION sub-types indicate the event either starts or ends at that location, respectively.

A PARTICIPANT may be a person, organization, or physical object. There are 3 PARTICIPANT sub-types: PRESENT, REMOTE, and TRANSITIONAL. The PRESENT sub-type indicates the participant is physically present at the location of the event. The REMOTE sub-type indicates that the participant is not physically present at the location of the event, yet participates nonetheless. This is possible largely through figurative language such as metonymy or metaphor (e.g., "*the United States entered the war*"). The TRANSITIONAL sub-type indicates that the participant was both present and remote, but at different times. This is common for motion events, where a participant may start within the spatial bounds of the event but finish elsewhere, or vice versa. For example, in the following sentence, *soldier* is present only at the beginning of the event, while *ballot* is the event's theme and thus present for the entire motion):

(3) The [soldier]$_{\text{PARTICIPANT}}$ [sent]$_{\text{E7}}$ his [ballot]$_{\text{PARTICIPANT}}$ for the [Maine]$_{\text{LOCATION}}$ [election]$_{\text{E8}}$.

The TRANSITIONAL sub-type itself has two further sub-types: P$\rightarrow$R (present to remote) and R$\rightarrow$P (remote to present).

3.3. Annotation Guidelines

Due to the fact that spatial containment relations between events are largely implicit, they often require some degree of intuition about the spatial bounds of an event. Beyond being exposed to a limited number of examples, an annotator must largely rely on his or her intuition about the spatial bounds of an event. This is not entirely unprecedented in natural language annotation, as Pan et al. (2011) asked annotators to provide their intuition for an event's temporal duration. For an example of how the annotator is asked to provide their intuition of an event's spatial bounds, given the text "*the [bombing] victim [died]*", the annotator must determine the expected spatial bounds for the *bombing* and *died* events, then determine if there is a relationship. This is highly intuitive (does one die immediately from a bombing, or does one's location change first?). Obviously, the context may aid in this tremendously. If the text above is followed by "*in the hospital*", one can reasonably assume the victim died in a different location, so the events would have a DIFFERENT relation (this relation is used as a negation of the four primary event relations when there is sufficient information to understand the events have non-intersecting spatial bounds). However, if the text above is followed by "*instantly*", one can reasonably assume the victim died on the scene of the bombing. In this case the most appropriate relation is to say that the *bombing* event CONTAINS the *died* event, since the spatial bounds for the bombing was likely larger.

Annotating spatial containment relations is further complicated by the fact that a document with n events has a possible $O(n^2)$ number of event relations. We therefore pro-

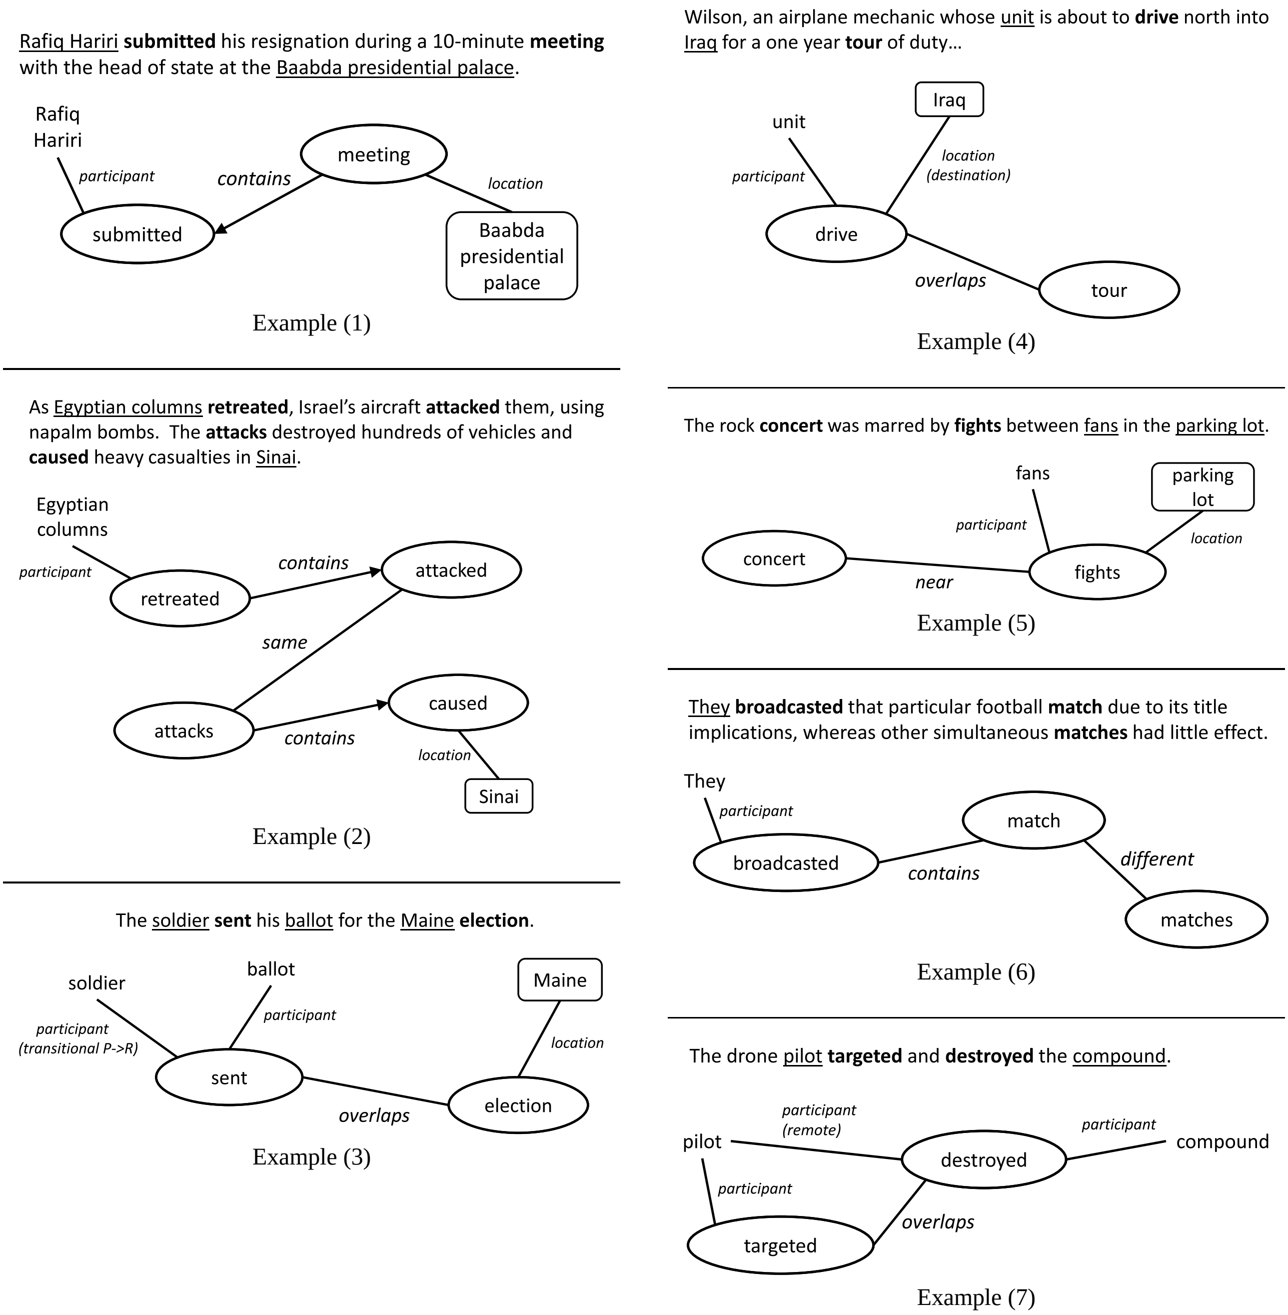

Figure 3: Spatial containment relations for Examples (1)-(7).

vided annotators with guidelines to simplify the annotation process[1]. For example, annotators were only required to look at the previous three sentences when searching for related events. Since most long-distance relations will be the result of event coreference, we feel this is a reasonable limitation. Furthermore, since spatial containment relations are transitive (i.e., if event A contains event B and B contains event C, then A contains C), there is no need to annotate the transitivity, as this can be computed automatically.

3.4. Additional Examples

Here we present examples to cover each type of relation. The three examples above as well as the following examples below are illustrated in Figure 3.

As previously discussed, in Example (1) the inference that *Rafiq Hariri* was in the *Baabda presidential palace* for event E1 (*submitted*) can be drawn by connecting E2 (*meeting*) to E1 with a CONTAINS relation. Similarly, in Example (2) the inference that *Egyptian columns* are spatially related to *Sinai* can be drawn with three separate relations, as shown in Figure 3. Note that a strict interpretation of the transitive closure of this graph does not place the *Egyptian columns* in *Sinai*, but only spatially related to it. We discuss this limitation as well as others in the next section. Example (3) illustrates a TRANSITIONAL PARTICIPANT (*soldier*), who is the agent of a motion event. The motion event E7 (*sent*) has an OVERLAPS relation with event E8 (*election*). Note that both Example (2) and (3) exhibit motion events (*retreated* and *sent*) intersecting with non-

[1]For more information, see our annotation guideline: http://www.hlt.utdallas.edu/∼kirk/spatial_containment_standard.pdf

motion events (*attacked* and *election*). The notable difference between these examples and why they merit different relations is the spatial bounds of the retreat is assumed to entirely encompass the spatial bounds of the attack (thus CONTAINS), while the send event does not encompass the entire election, assumed to take place across the entire state of Maine (thus OVERLAPS).

Example (4) also exhibits an OVERLAPS relation:

(4) Wilson, an airplane mechanic whose [unit]$_{\text{PARTICIPANT}}$ is about to [drive]$_{\text{E9}}$ north into [Iraq]$_{\text{LOCATION}}$ for a one year [tour]$_{\text{E10}}$ of duty, put his finger on a problem that has bedeviled the Pentagon for more than a year.

Here, it is assumed that event E9 (*drive*) will only cover a subset of the spatial bounds of event E10 (*tour*). This example also exhibits a non-COMPLETE LOCATION. Since event E7 starts from outside *Iraq*, the LOCATION is marked as the DESTINATION.

Example (5) exhibits a NEAR relation:

(5) The rock [concert]$_{\text{E11}}$ was marred by [fights]$_{\text{E12}}$ between [fans]$_{\text{PARTICIPANT}}$ in the [parking lot]$_{\text{LOCATION}}$.

Here, the spatial bounds of event E11 (*concert*) are assumed to be limited to the stage and audience area (e.g., a concert hall or field). Event E12 (*fights*) is instead related with a NEAR relation. This allows us to rule out the incorrect inference that the *concert* was located in the *parking lot*.

Example (6) exhibits a DIFFERENT relation:

(6) [They]$_{\text{PARTICIPANT}}$ [broadcasted]$_{\text{E11}}$ that particular football [match]$_{\text{E14}}$ due to its title implications, whereas other simultaneous [matches]$_{\text{E15}}$ had little effect.

Here, event E14 (*match*) is marked as DIFFERENT from event E15 (*matches*). This example demonstrates both our primary motivations for including the DIFFERENT relation: (1) to provide examples where possibly coreferential events are both non-coreferential and do not share spatial bounds, and (2) to explicitly state that two events are not spatially related. The spatial bounds of two events are not necessarily DIFFERENT if they cannot be connected by a transitive closure operation (instead, their spatial relationship is simply unknown), so the DIFFERENT relation allows for such an explicit statement when clear.

Example (7) exhibits a PARTICIPANT classified as REMOTE:

(7) The drone [pilot]$_{\text{PARTICIPANT}}$ [targeted]$_{\text{E16}}$ and [destroyed]$_{\text{E17}}$ the [compound]$_{\text{PARTICIPANT}}$.

Since the *pilot* is not present at the *compound*'s destruction (instead, the *drone* is present), he or she is considered REMOTE for event E17 (*destroyed*) but PRESENT in event E16 (*targeted*). This is based on the intuition that the targeting is done both locally (i.e., by the pilot and his/her control center) and remotely (i.e., by the drone), while the destroying is done entirely by the drone and it's weapons systems. The events are connected via an OVERLAPS relation as the *pilot* is not considered spatially part of E17 (*destroyed*), nor is the *compound* considered spatially part of E16 (*targeted*), so neither is a subset of the other.

Note that in many of these examples there exists the possibility of multiple valid interpretations. While this certainly makes annotation difficult, it also highlights the vague nature of implicit spatial relations and reinforces our decision for a relatively simple set of spatial relations.

3.5. Current Limitations

We believe the primary limitation of our relations center around the lack of granularity in the event, PARTICIPANT, and LOCATION relations. The choice of a level of granularity plays an important role in natural language inference (Hobbs, 1985a). When deciding the proper level of detail for our representation, we took a pragmatic approach based on two competing factors: (1) increasing granularity raises the level of difficulty of annotation, lowering annotator agreement on an already difficult annotation task and thus reducing the effectiveness of an approach based on our data, and (2) decreasing granularity reduces inferential power, reducing the effectiveness of an approach based on our data. We therefore chose to use a basic set of relations, with a few key exceptions as highlighted above. If automated methods prove successful on this data, the issue of granularity may be re-visited and the annotation revised for many of the following limitations.

The examples previously discussed give an indicator of the strengths and limitations of our set of relations, particularly as they relate to motion events. We integrated basic support for motion end-points into PARTICIPANTs (through the TRANSITIONAL sub-type) and LOCATIONs (through the SOURCE and DESTINATION sub-types), but omitted many of the other properties of motion from these relations. Further, we omitted motion properties from the event relations entirely.

Example (4) illustrates how this lack of granularity may transfer into sub-optimal inferences. We would like to know that event E10 (*tour*) takes place in *Iraq*, but that LOCATION is only attached to event E9 (*drive*) as a DESTINATION. Using two relations instead of one would allow for the desired inference: (1) at the beginning of the *drive* motion event, the relation between E9 and E10 would be DIFFERENT (or NEAR if indicated by the context), (2) at the end of the *drive* motion event, the relation would be that E10 CONTAINS E9. Since *Iraq* is a DESTINATION LOCATION, we could then infer that at least part of the *tour* was located in *Iraq*. Similar inferences could be made for Example (3) to determine that the *ballot*'s final LOCATION was in *Maine*.

Other motion properties could be useful as well, such as the path of a motion event. Instead of simply representing the SOURCE and DESTINATION, integrating arbitrary midpoints and directions (possibly including a temporal component) would allow for additional inferences to be made. In natural language text, however, such details about motion events are usually omitted as the spatial properties of events are almost always underspecified.

Other granularity limitations involve PARTICIPANT and LOCATION relations. Some PARTICIPANTs, such as the *Egyptian columns* in Example (2), largely define the spatial bounds of the event (e.g., the spatial bounds of the retreat

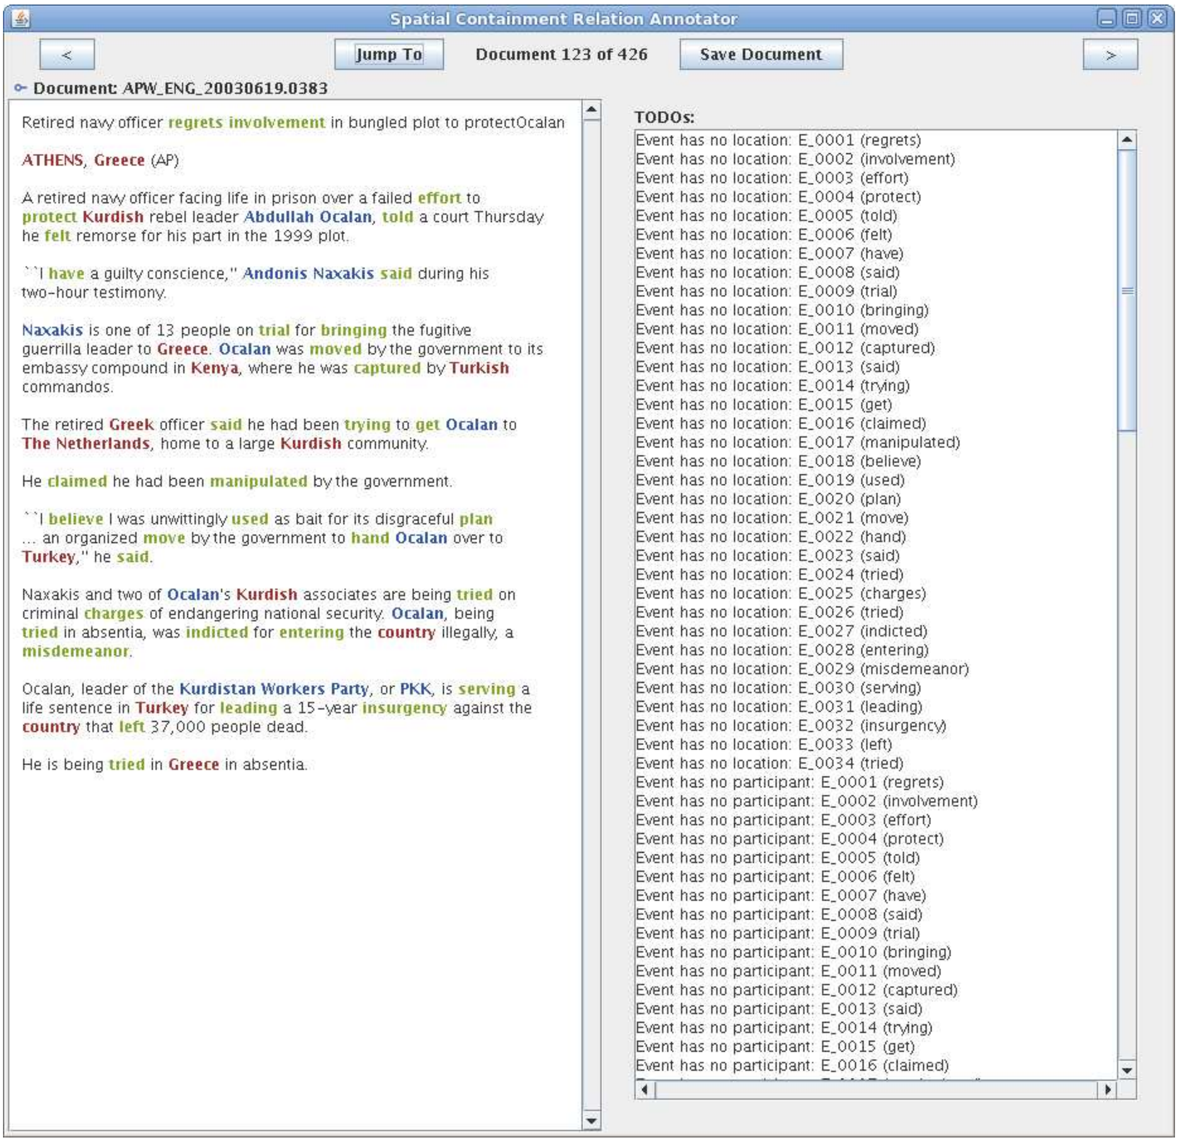

Figure 4: Spatial containment relation annotator.

were defined by the locations of the retreating columns), while other PARTICIPANTs form a small part of the overall event (e.g., the *ballot* in Example (3) is a small part of the *election* event). Similarly, knowing how the LOCATION is spatially related to the event would be useful for inference. For now, we simply assume the event CONTAINS the PARTICIPANT, while the LOCATION CONTAINS the event. As with all the previously mentioned granularity issues, our primary goal with this data is to create a set of basic event relations conveying implicit spatial containment relationships. If, in the future, automated methods can achieve sufficient accuracy on these basic relations, the granularity of our annotations can be increased to suit inference needs.

4. Corpus Creation and Analysis

We chose to annotate our spatial containment relations on the SpatialML corpus (Mani et al., 2008). SpatialML already contains a wealth of spatial information, including location mentions, their gazetteer normalizations, and relations between locations. Thus it is natural to use our spatial event relations to augment the corpus's existing spatial information. Many of the SpatialML documents are conversations and broadcast transcripts, which we do not expect to contain a significant amount of spatial event information, so our annotation effort has focused on annotating the 160 documents in SpatialML derived from newswire.

Annotators were provided with a custom annotation tool, shown in Figure 4, for efficient annotation. This tool simplifies the process of searching for related events and enforces consistency in the annotations. To aid the annotators, we automatically annotated events using TARSQI (Saurí et al., 2005) and person/organization entities using BIOS[2].

To get a feel for the difficulty of the task, we gave a brief overview of the task to our three annotators (without showing them the annotation guideline) and asked them to each annotate the same five documents. As expected, initial agreement was very low when evaluated with Fleiss' Kappa (Fleiss, 1971). Initial agreement on whether two events are related (without the relation type) was 0.23, which falls into the "fair agreement" category from Lan-

[2] http://www.surdeanu.name/mihai/bios/

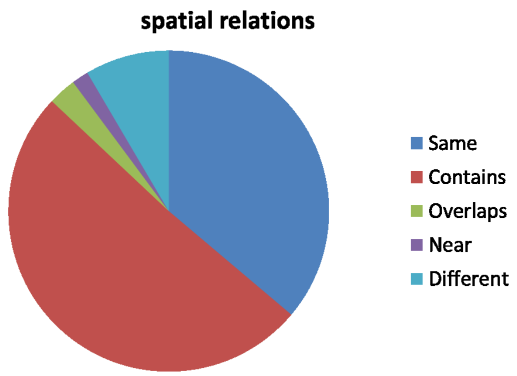

Figure 5: Distribution of spatial containment relations in our corpus.

dis and Koch (1977). Agreement on the relation type was 0.51 ("moderate agreement"). After providing the annotation guideline and reviewing disagreements, the annotators proceeded to individually annotate approximately 12 documents each (depending upon length). On average, a SpatialML newswire document took an annotator approximately one hour to annotate. Next, three documents were chosen to be annotated by all annotators, showing improved agreement. Agreement on whether two events are related for these documents was 0.45 ("moderate agreement"), while agreement on the relation type improved to 0.64 ("substantial agreement"). Because these relations are largely implicit and based entirely on the annotator's spatial interpretation of the event, it is likely that near-perfect agreement is not a practical goal. After this, we have proceeded with single-annotation for the remaining documents. Currently, approximately half of the newswire documents from SpatialML have been annotated by at least one annotator. We plan to make an initial version of these annotations publicly available soon.

Figure 5 illustrates the distribution of spatial containment relations in our annotated corpus. Clearly, SAME and CONTAINS stand out as the most annotated relations, while NEAR and OVERLAPS are rare. This is largely due to the fact that these relations require more contextual justification than the SAME and CONTAINS relations. SAME and CONTAINS can usually be determined through an intuitive understanding of each event's semantics.

To evaluate the difficulty of this task, we propose a few simple, supervised features to act as a baseline and illustrate the importance of integrating more semantic components. The four features are:

1. The two event words (e.g., the feature value for the first event pair from Example (1) would be `submitted::meeting`).

2. The two event lemmas (e.g., `submit::meeting`).

3. The words between the events (e.g., `his`, `resignation`, `during`, `a`, `10-minute`).

4. The hypernyms of the events using a first-sense assumption (e.g., `refer::gathering`, `send::gathering`, ..., `move::abstraction`, `move:entity`).

Relation Exists (binary)			
Feature Set	P	R	F_1
EW	13.7	31.0	19.0
EL	14.7	33.4	20.4
WB	14.3	57.9	22.9
HN	17.9	32.6	23.1
EW + HN	19.7	31.4	24.2
EL + HN	19.4	31.5	24.0
WB + HN	25.8	39.9	31.3
EW + WB + HN	27.8	37.5	31.9
EL + WB + HN	27.6	38.0	32.0
EW + EL + WB + HN	29.1	35.5	32.0

Table 1: Baseline experiments for whether two spatial events are connected by a spatial containment relation. EW = event words; EL = event lemmas; WB = words between, HN = hypernyms.

Relation Type (5-way)	
Feature Set	%
EW	58.3
EL	57.7
WB	52.1
HN	54.9
EL + EW	57.9
WB + EW	53.1
HN + EW	54.8

Table 2: Baseline experiments for the type of spatial containment relation that connects two events. See Table 1 for legend.

We experimented with these features using a support vector machine implementation (Fan et al., 2008).

The results of these experiments are shown in Tables 1 and 2. These results indicate that while automatic recognition of spatial containment relations between events is possible, a far richer set of semantic features is necessary to both automatically recognize and categorize these relations. Given that only around 10% of event pairs within a 3-sentence window are marked as having a spatial containment relation, an F_1-measure of 32.0 shows that even basic lexico-semantic methods can capture many cases of spatial containment relations.

The best-performing experiment for relation type classification uses only the words in the two events. Neither the context words between the events or the use of hypernyms to generalize the events improve relation type classification. However, the result using this feature is still quite poor given that the most frequent class baseline is around 50%. This suggests that the features in these baseline experiments do not capture the relevant spatial information. Given that the WB (word between) feature performed so poorly, it is likely that the lexical context offers little evidence of the relation type. This validates our assertion that this task is largely implicit and requires some combination of discourse clues and world knowledge about the semantics of the two related events.

Upon analysis, it seems clear that in order for an automatic method to prove succesful on this task, it must incorporate: (1) an understanding of event semantics to represent how pairs of events are related (such as using event scenarios (Bejan, 2008) or narrative schemas (Chambers and Jurafsky, 2008)), (2) event coreference to form chains of identical events, (3) discourse relations that hold between events, and (4) a sense of the relative spatial bounds of events (e.g., events that happen at the level of cities and nations as opposed to those that happen at the level of individual human interactions).

5. Conclusion

In this paper, we have discussed our motivation and annotation schema for spatial containment relations between events, placing it within previous work in both event relations (e.g., TimeML) and spatial representation (e.g., RCC-8, SpatialML). We described our process for creating a corpus with these event relations and analyzed the current state of our corpus, which is still undergoing development. We performed a set of baseline experiments with simple lexico-semantic features in order to determine the feasibility of using these annotations for creating an automatic system for detecting spatial containment relations between events. In our analysis, we outlined several key components necessary to perform automatic recognition of our event relations, including event semantics, event coreference, discourse relations, and approximate spatial bounding. For future work, we plan to integrate many of these approaches into an automatic approach for recognizing the spatial containment relations between events.

6. References

Cosmin Adrian Bejan and Sanda Harabagiu. 2010. Unsupervised Event Coreference Resolution with Rich Linguistic Features. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*.

Cosmin Adrian Bejan. 2008. Unsupervised Discovery of Event Scenarios from Texts. In *Proceedings of the 21st Florida Artificial Intelligence Research Society International Conference (FLAIRS)*.

Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised Learning of Narrative Schemas and their Participants. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*.

Zheng Chen, Heng Ji, and Robert Haralick. 2009. A pairwise event coreference model, feature impact and evaluation for event coreference resolution. In *Proceedings of the RANLP Workshop on Events in Emerging Text Types*, pages 17–22.

Quang Xuan Do, Yee Seng Chan, and Dan Roth. 2011. Minimally Supervised Event Causality Identification. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303.

Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A Library for Large Linear Classification. *Journal of Machine Learning Research*, 9:1871–1874.

Joseph L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382.

Jerry R. Hobbs. 1985a. Granularity. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 432–435.

Jerry R. Hobbs. 1985b. On the Coherence and Structure of Discourse. Technical Report CLSI-85-7, Center for the Study of Language and Information.

J. Richard Landis and Gary G. Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33:159–174.

Inderjeet Mani, Janet Hitzeman, Justin Richer, Dave Harris, Rob Quimby, and Ben Wellner. 2008. SpatialML: Annotation Scheme, Corpora, and Tools. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation*.

Martha S. Palmer, Deborah A. Dahl, Rebecca J. Schiffman, Lynette Hirschman, Marcia Linebarger, and John Dowding. 1986. Recovering implicit information. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pages 10–19.

Feng Pan, Rutu Mulkar-Mehta, and Jerry R. Hobbs. 2011. Annotating and Learning Event Durations in Text. *Computational Linguistics*, 37(4):727–752.

James Pustejovsky and Jessica L. Moszkowicz. 2008. Integrating Motion Predicate Classes with Spatial and Temporal Annotations. In *Proceedings of COLING 2008*, pages 95–98.

James Pustejovsky, José Castano, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir Radev. 2003a. TimeML: Robust Specification of Event and Temporal Expressions in Text. In *IWCS-5 Fifth International Workshop on Computational Semantics*.

James Pustejovsky, Patrick Hanks, Roser Saurí, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, and Marcia Lazo. 2003b. The TIMEBANK Corpus. In *Proceedings of Corpus Linguistics*, pages 647–656.

James Pustejovsky, Jessica L. Moszkowicz, and Marc Verhagen. 2011. Iso-space: The annotation of spatial information in language. In *Proceedings of the Sixth Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation*, pages 1–9.

David A. Randell, Zhan Cui, and Anthony G. Cohn. 1992. A Spatial Logic based on Regions and Connection. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, volume 117.

Roser Saurí, Robert Knippen, Marc Verhagen, and James Pustejovsky. 2005. Evita: A Robust Event Recognizer For QA Systems. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT-EMNLP)*, pages 700–707.