

Inferring Implicit Event Locations from Context with Distributional Similarities

Jin-Woo Chung, Wonsuk Yang, Jinseon You, Jong C. Park*

School of Computing, KAIST, Republic of Korea
 {jwchung, derrick0511, jsyou, park}@nlp.kaist.ac.kr

Abstract

Automatic event location extraction from text plays a crucial role in many applications such as infectious disease surveillance and natural disaster monitoring. The fundamental limitation of previous work such as SpaceEval is the limited scope of extraction, targeting only at locations that are explicitly stated in a syntactic structure. This leads to missing a lot of implicit information inferable from context in a document, which amounts to nearly 40% of the entire location information. To overcome this limitation for the first time, we present a system that infers the implicit event locations from a given document. Our system exploits distributional semantics, based on the hypothesis that if two events are described by similar expressions, it is likely that they occur in the same location. For example, if “*A bomb exploded causing 30 victims*” and “*many people died from terrorist attack in Boston*” are reported in the same document, it is highly likely that the bomb exploded in Boston. Our system shows good performance of a 0.58 F1-score, where state-of-the-art classifiers for intra-sentential spatiotemporal relations achieve around 0.60 F1-scores.

1 Introduction

Natural language processing on plain documents has enabled real-time, elaborate monitoring for the occurrence of infectious disease and natural disasters [Barboza *et al.*, 2014]. For example, when the postings such as “*I’m not feeling well for 3 days, but I don’t know why*” are suddenly found at the same time with their common location information, it is likely that an infectious disease has been spread out in the location [Hartley *et al.*, 2013]. Moreover, the same logic has been utilized for natural disaster monitoring using Twitter postings [Imran *et al.*, 2016]. The location inference from text has its unique advantage compared to GPS based approaches with respect to privacy issues since such information is usually reported by its writers voluntarily.

Major League Baseball plans to warn its players heading to [Canada, particularly to Toronto LOCATION] for the All-star break in July to [take EVENT] some precautions. That means instead of [meeting EVENT] the fans and [signing EVENT] autographs, they’ll be [keeping EVENT] their distance and [waving EVENT].

Figure 1: Events and their location expressed in a document

However, the inference of such location information is a challenging task because general postings do not necessarily report it in a regularized form.

In general, when people write an article about a series of real-world events, they do not specify location information for every event as it is naturally understood by context. Instead, people usually mention location information only once for a certain event, and then omit it for other events if they occur in the same location. For example, one can easily understand that all the events marked in Figure 1 occur (or is expected to occur) in a place expressed by “*Canada, particularly to Toronto*”.

This means that for many events in a document, location information could be implicit and not directly stated in their local syntactic structure, and that there should be some degree of inference to find a proper location from context. However, most work on extracting spatial information from text has focused only on locations that are explicitly mentioned together with events in a sentence and left this essential inference problem as future work.

Motivated by the importance of this issue, recent work has introduced detailed analysis and manual annotations of location information of events that is not directly stated but still inferable from context in a document [Chung *et al.*, 2015]. It shows that about half of the entire event-location pairs in news articles are not directly stated via explicit locative argument structure, and that nearly 40% of all the pairs can be inferred from different sentences in a document. This highlights the need of inference for event location extraction from text.

To the best of our knowledge, no inference system has been developed for the purpose of extracting implicit location information from text. The most relevant and recent attempts are made through the SpaceEval shared task

* Corresponding author

[Pustejovsky *et al.*, 2015], especially focusing on extracting MOVELINKS, but they deal mainly with classifying the relations between motion events and their arguments that are directly stated together in the same sentence. Here, we present an inference system that exploits distributional semantics to spread explicitly stated location information to other events in a document whose location is implicit. For the extraction of the explicitly stated locations, we exploit several linguistic features inspired by the state-of-the-art systems in SpaceEval to meet our purpose.

The key intuition behind using distributional semantics for implicit location extraction is that if two events mentioned in a coherent document are described by expressions showing high distributional similarities, it is likely that they occur in the same location. For example, if two sentences “*A bomb exploded causing 30 victims*” and “*many people died from terrorist attack in Boston*” are mentioned in the same document, it is highly likely that the *bomb explosion* also occurs in *Boston*. Based on this idea, we apply distributional semantic models to measure the similarities between words composing two events. We also evaluate our methods on the annotated corpus, achieving an F1-score of 0.62 for all locations and 0.53 for implicit locations only. Considering the highly implicit nature of the document-level location inference task, our system shows great initial results for this new task.

2 Related Work

Research on extracting location information of events has received attention more recently than work on other event properties, such as temporal relations, causal relations, and coreference. There have been efforts into establishing annotation languages and specifications for describing expressions referring to geographic regions, such as SpatialML, which is inspired by TimeML [Pustejovsky *et al.*, 2003a], and Spatiotemporal Markup Language [Mani *et al.*, 2008], both of which have been integrated into ISO-Space [Pustejovsky *et al.*, 2011]. Another research line is Spatial Role Labeling [Kordjamshidi *et al.*, 2010], which addresses extracting spatial relations between objects triggered by spatial indicators such as *in*, *on*, and *under*, and assigning specific roles to objects with respect to the triggering indicators. It has been introduced in the subtasks of SemEval 2012 and 2013 [Kordjamshidi *et al.*, 2012; Kolomiyets *et al.*, 2013], and has recently been extended to SpaceEval [Pustejovsky *et al.*, 2015], which also incorporates specifications from ISO-Space to enrich the granularity of the spatial semantics covered by the previous tasks. SpaceEval also highlights motion events and their arguments under the concepts of dynamic spatial relations, and considers extracting spatial relations among objects in motion, or MOVELINKS, as a key challenge. However, these lines of research have focused mainly on employing fine-grained spatial concepts to identify complex relations and triggering signals that are stated within a sentence, but not on recognizing implicit, non-stated relations between general non-motion events and locations from context in a document.

Vempala and Blanco (2016) address the problem of inferring temporally-anchored spatial knowledge from PropBank-style semantic role representations. Their specific task is to determine whether entities of a given event are located in one of its arguments found in the same sentence, and to temporally anchor their spatial relationship with respect to the event. Although their work deals with inferring implicit spatiotemporal knowledge, its scope is limited to verb-argument relations within single sentences.

Another line of work is the recognition of spatial containment relations between events [Roberts *et al.*, 2012; Roberts *et al.*, 2013]. They aim at determining whether the spatial boundaries of two events overlap or not, as exemplified in “*The presentation at the conference was excellent*”, where the *presentation* event is spatially contained by the *conference* event. Although they limit the scope of relations only to a 3-sentence window and do not directly recognize location expressions for given events, they discuss important implication for spatial relations: If two events occur together in narratives frequently, i.e., with high statistical associations, it is likely that they are spatially related. While they rely on simple co-occurrence based statistics to measure the associations between event words, we explore the use of distributional similarities, which have recently been shown to achieve good generalization capabilities in many natural language processing tasks.

3 Data and Annotations

For experiments and evaluation, we use the corpus presented in Chung *et al.* (2015), which, to the best of our knowledge, is the only work that provides manual annotations of event-location relations on a document level. We briefly describe the corpus in Section 3.1 and present newly added annotations of containment and coreference relations among location expressions in Section 3.2 for more practical evaluation.

3.1 Corpus

The corpus provides annotations of event mentions and their location expressions (i.e., event-location links) as found within a given document, which best indicate where those events happen. They are annotated on top of 48 documents from CNN broadcast news articles in OntoNotes Release 5.0 [Hovy *et al.*, 2006], covering diverse topics such as social issues, accidents, wars, politics, and sports. An event mention is any single word token corresponding to verb-category words including base verbs, inflected verbs, gerunds, but excluding non-content words such as copula and auxiliary verbs, as adopted in many event-related corpora such as TimeBank [Pustejovsky *et al.*, 2003b]. Note, however, that the corpus does not consider noun phrases as event mentions, which is also the limitation of the present work. A location expression is any word or sequence of words that indicates where given events occur. Each event mention is linked to at most two location expressions. Every location expression is linked to at least one event mention. The corpus has 846 event mentions, 337 locations expressions, and 721

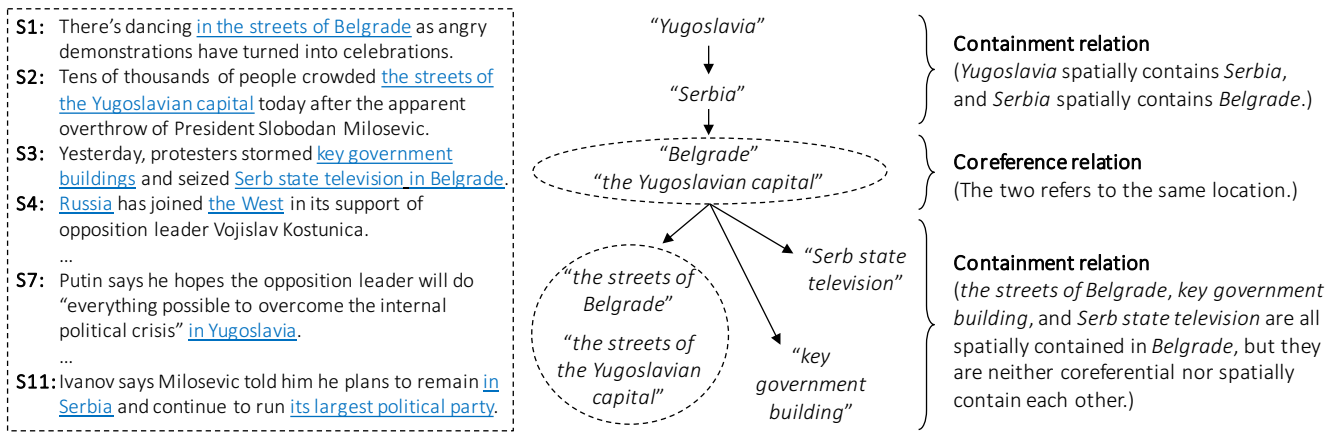


Figure 3: Example coreference and containment relations between locations in a news article

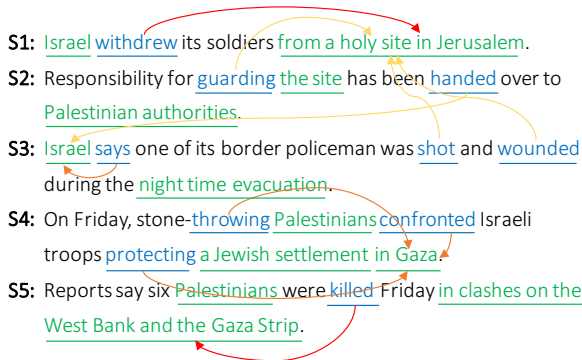


Figure 2: Example of event-location links. Expressions in blue and green are event mentions and location expressions, respectively. Red links connect events and their directly stated locations, i.e., explicit locative arguments. Orange links are created by additional intra-sentential cues. Yellow links are inter-sentential links.

event-location links in total. Figure 2 shows an example of event-location links.

The OntoNotes corpus: The annotations in the corpus we use for our experiments is created on top of the OntoNotes 5.0 corpus [Hovy *et al.*, 2006], which provides several layers of manually annotated linguistic information such as part-of-speech tags, parse trees, semantic role labels, word sense, named entities, and coreference. In this work, we use gold information provided in these OntoNotes annotations for consistent analysis of contributions of such features to overall performance.

3.2 Annotating Relations among Locations

Although the corpus described above provides a good basis for our experiment, it lacks two kinds of information crucial to the proper evaluation of the event-location relations: coreference and containment relations among location expressions. This is inspired by complex, hierarchical nature of relations between three-dimensional space. Figure 3 shows an example of these relations existing in the corpus. Here, *Yugoslavia* spatially contains *Serbia*, which again spatially contains *Belgrade*. Moreover, two coreferential expressions, *Belgrade* and *the Yugoslavian capital*, refer to

the same location. This means that for a given event mention, more than one location expression could be correct answers. For example, let us assume that a certain event is linked to *Belgrade* in the annotated corpus, but a system chooses *the Yugoslavian capital* as an answer. Here, it can be said that the system's answer is still correct as both *Belgrade* and *the Yugoslavian capital* refer to the same location. Moreover, if the system chooses *Serbia* instead, it can also be considered another correct answer in a broad sense (although it is not the most precise) as it still spatially contains *Belgrade*. However, if the system chooses *key government building*, it must not be a correct answer as it is only a spatial subset of *Belgrade*. Since the original corpus does not include information of this kind (coreference and containment), which may prevent fair evaluation, we performed additional annotations, according to the procedure below.

Annotation procedure and result

Two annotators participated in the manual annotation where they were asked to search given documents for a possible relation between two arbitrary annotated location expressions, as explained above, considering the context. We did not ask them to annotate redundant relations that already hold by transitivity; i.e., if they annotate two relations "A contains B" and "B contains C", they do not annotate "A contains C" as it can be naturally inferred by the two relations. Through this annotation, we obtained 85 coreference relations (55 coreference chains) and 172 containment relations (i.e., location pairs) in total. To calculate the inter-annotator agreement, we used bidirectional precision/recall-based measure as adopted in discourse relation research (e.g., [Mírovský *et al.*, 2010])¹. We

¹ We did not use conventional measures for inter-annotator agreement such as Cohen's kappa and Dice's coefficient for two reasons: (1) Annotators must decide first what to annotate by choosing two arbitrary locations, before choosing a label. (2) The agreement may not be symmetric; e.g., if annotator 1 annotates "X contains Y" & "Y contains Z", whereas annotator 2 "X contains Z" only, we can say that annotator 1 agrees with annotator 2 by transitivity, but not vice versa, because annotator 2 misses Y.

obtained an agreement of 0.79, which is quite reasonable, given that this type of document-level relation annotation usually involves a heavy cognitive burden, as discussed in [Roberts *et al.*, 2012].

4 Method

Our system for extracting event-location links consists of two steps: detection of explicitly stated locations and detection of implicit locations. In the first step, we recognize candidate expressions that are eligible for locations first, and then detect event-location links that are explicitly stated in a sentence, i.e. those that are signaled by locative semantic roles in the OntoNotes corpus. In the second step, we detect other intra-sentential links based on additional cues on a sentence level, and then use them to infer inter-sentential links based on coreference relations and distributional similarities. Figure 2 shows how event-location links are recognized through the steps from location expression extraction to inter-sentential link recognition.

4.1 Detection of Explicitly Stated Locations

Although we do not restrict location expressions to conventional locative arguments, it would not make the process feasible if we consider all possible sequences of words as candidates for location expressions. We thus consider the following types of expression as initial candidates, which we believe cover most, though not all, location expressions: (1) noun phrases and adverbial phrases, or (2) adjectival forms of place names and their demonymic equivalents, e.g., *Russia* and *Russian*. To extract them, we use phrase bracketing tags from parse trees and place/country names from Wikipedia pages². We also filter out candidates with the following OntoNotes named entity tags that indicate non-location entities: PERSON, LANGUAGE, DATE, TIME, PERCENT, MONEY, QUANTITY, ORDINAL, and CARDINAL. Our system traverses the parse tree of each sentence and collects candidates satisfying these conditions, obtaining 1649 candidate expressions in total. While we found that they cover 98% of the whole gold location expressions, they also generate too many incorrect ones, considering that there are only 337 gold expressions. However, we did not filter them out any further since it may significantly lower the recall and may also limit the diversity of location information.

To recognize links between events and their explicitly stated locations, the system finds expressions that act as a locative semantic argument of given events, which is indicated by two locative semantic role labels in the OntoNotes annotations: ARGM-LOC and ARGM-DIR. Note that these labels are, by their definition, expected to strongly signal locations for given events and are also used in many systems participating in SpaceEval [Pustejovsky *et al.*, 2015]. Example expressions are “*from a holy site in Jerusalem*” and “*in clashes on the West Bank and the Gaza Strip*” in Figure 2

² https://en.wikipedia.org/wiki/List_of_adjectival_and_demony_mic_forms_of_place_names
https://en.wikipedia.org/wiki/List_of_adjectival_and_demony_mic_forms_for_countries_and_nations

which are connected to events *withdraw* and *kill*, respectively, by red links. Note that only two of the ten links in Figure 2 can be recognized by explicit locative argument structure.

4.2 Detection of Implicit Locations

Recognition of additional intra-sentential relations

We note that there are some potential links between events and locations that are mentioned in the same sentence but are not connected via locative semantic role labels, which we cover by the following three factors. We create a link for events when their locative semantic argument is missing (i.e., their location is not directly stated) but there is another location in the same sentence satisfying one of the three conditions.

Factor 1) Locative semantic argument of another event that shares the same first argument: This is based on the assumption that if the same agent is involved in two consecutive events in a sentence and a location is explicitly stated for only one of them, it is likely that the other event also occurs in that location, as shown in example (1).

Factor 2) Locative prepositional phrases modifying the first or second argument of a given event: In some cases, locations are syntactically attached to the first or second argument of events (i.e., their subject or object) as an adverbial phrase with locative prepositions rather than used as their explicit locative argument, as shown in example (2).

Factor 3) Nominative country/city/organization: When the event agent (nominative entity) refers to a large community such as countries, cities, and organizations, an actual event often takes effect by its members involved and in this case the event agent itself indicates locations, as shown in example (3).

- (1) Annan also met with Arafat [in Gaza LOC] and [offered EVENT] proposals for ending the crisis.
- (2) Students [at a middle school in Calaveras County, California LOC], are [getting EVENT] an unwanted lesson in entomology.
- (3) [Canada LOC] is [saying EVENT] good-bye to a former leader.

We show that these cues above work quite effectively for recognizing additional intra-sentential links, helping to achieve high precision, to be discussed in Section 5.

Recognition of inter-sentential relations

We utilize explicitly stated locations and implicit intra-sentential links to infer inter-sentential links on a document level. We assume that for events whose locations are not directly stated, it is possible to infer them from the context of a given document, using other event-location links that are already recognized. In other words, the inference is done in a way that location information that is directly stated (and is thus already recognized) spreads out across neighboring sentences.

We first utilize coreference relations of OntoNotes between the first argument in a similar way to Factor 1 of additional intra-sentential links, by extending its assumption

	Relaxed match			Strict match		
	P	R	F	P	R	F
Intra-sentential only	81.1	27.6	41.2	73.7	25.0	37.4
Inter-sentential (all with coreference)	P	R	F	P	R	F
Nothing (only coreference)	73.5	37.7	49.9	64.8	33.2	43.9
Method 1 (event-event distributional similarity)	70.6	39.0	50.3	62.0	34.3	44.1
Method 2 (event-location distributional similarity)	43.1	43.1	43.1	37.9	37.9	37.9
Method 3 (first argument distributional similarity)	55.7	52.3	54.0	39.8	37.4	38.6
Method 4 (maximum)	66.5	47.2	55.2	53.0	37.7	44.1
Method 5 (average)	66.7	47.4	55.4	53.7	38.2	44.6
Method 4 + Coarse-grained	62.1	62.1	62.1	43.7	43.7	43.7
Method 5 + Coarse-grained	62.3	62.3	62.3	44.1	44.1	44.1
Method 5 + Coarse-grained (when evaluated for implicit locations only)	58.8	58.8	58.8	39.2	39.2	39.2

Table 1: Experimental results

to the inter-sentential scope. For instance, example (4) shows two consecutive sentences, where two events *stay* and *spoke* have the coreferential arguments (i.e., *he*) and the system thus creates an inter-sentential link between *spoke* and *in Washington*.

(4) President Clinton has cancelled all of his plans for the day so [he COREF] can [stay EVENT] [in Washington LOC] and monitor developments in the Middle East. This morning [he COREF] [spoke EVENT] to Ehud Barak and Yasser Arafat.

We then apply distributional similarities as our key contextual features and test how well they work for such inference. More specifically, for a given event (E1) whose location is not recognized, the system searches all the other events whose locations are already recognized, and finds the one (E2) that shares the highest distributional similarity. It then creates a link between E1 and the location linked to E2. For example, for events *shot* and *wounded* in Figure 2, if they share the highest distributional similarity with the *withdrew* event, among others, we link them to “*from a holy site in Jerusalem*”, an explicit locative argument of *withdrew*.

5 Experiments

5.1 Setting

To calculate the distributional similarity score, we use the *gensim* library³ and the *word2vec* model, which is one of the most popular distributional semantic models for words known as a type of word embedding [Mikolov *et al.*, 2013]. We use the *word2vec* representations pre-trained on part of Google News dataset that consists of 100 billion words⁴.

Configurations for the distributional similarities

We experiment with five different settings as shown below to see in what configuration the distributional similarities contribute to the overall performance. For example, Method 1 is based on the assumption that the similarities between event words (verbs) best determine whether they share the

same location or not, whereas Methods 4 and 5 are based on the assumption that event words, their first arguments, and location expressions together contribute to such decision.

- **Method 1)** Event-event (verb-verb) similarities
- **Method 2)** Event-location (verb-noun) similarities
- **Method 3)** First argument (noun-noun) similarities
- **Method 4)** Maximum values of Methods 1-3
- **Method 5)** Average values of Methods 1-3

Dealing with events that remain not linked

If some events still remain not linked to any location expression in a given document even after applying distributional similarities, we take two coarse-grained approaches to link them to one of the candidate locations. The first approach is to choose the location expression that refers to the biggest location in a given document. This is based on the intuition that, for certain types of events, it is quite difficult to choose a specific location, and that such events can be linked to a location expression that refers to a large, vague region such as countries and cities. Our system relies on a simplifying assumption such that continents are the biggest location information, followed by countries, cities, and organizations in the order of size. We use place/country names, and types of named entity tags as used in the step for candidate location expression extraction (Section 4.1), to roughly compare the size of space. The second approach is to choose the location expression linked to another event mention that is closest to the given event mention. The closeness here is measured by textual distance in a document. If the two approaches above produce different location expressions for the given event mention, we choose the one that shares a higher distributional similarity with the given event mention.

Measuring the performance

For each event mention, we compare the location expression predicted by our system against the gold one in our corpus. We consider two location expressions to match each other if they share the same head word (noun), because detecting the exact boundary of noun/adverbial phases is not the main focus of this work. The performance is measured by counting how many event mentions are correctly linked to their gold location expression. We use precision, recall, and F1-scores

³ <http://radimrehurek.com/gensim>

⁴ <https://code.google.com/archive/p/word2vec/>

to report final performance. We also apply two different measures, strict match and relaxed match, according to whether matches via containment relations (Section 3.2) are considered correct or not. Note that matches via coreference relations are always considered correct.

5.2 Results and Discussion

Experimental results

Table 1 shows experimental results with different settings. The system produces good results when it employs combinations (maximum or average) of different measures of distributional similarities and coarse-grained methods.

While the performance itself is not very high, it still seems to provide great initial results for this new task, given that general state-of-the-art performance of classifying spatiotemporal relations lies around F1-scores of 0.60; e.g., 0.62 for temporal relation type classification in TempEval-3 [Mirza and Tonelli, 2016] and 0.60 for the extraction of MOVELINKS in SpaceEval [D’Souza and Ng, 2015].

The results show that using the intra-sentential method which relies only on local information achieves high precision, but its low coverage leads to poorer overall performance than the inter-sentential methods that exploit coreference and distributional similarities. Moreover, it is shown that for measuring distributional similarities, we obtain much better results when considering maximum or average values than considering only single types of event element pairs (e.g., event-event or event-location). This suggests that spatial relatedness of two events is better determined by distributional properties of several event elements and target locations combined together.

Another interesting point is that the coarse-grained method significantly increases the performance in the relaxed setting, but does not produce the best result in the strict setting. There would be two plausible reasons for this. The first would be that plenty of events in news articles are difficult to unambiguously associate with some specific location, such as politics-related events. The other would be that it might be safer to choose the expression referring to large space for given events unless the system has some degree of confidence in giving correct answers.

Error analysis and possible improvements

We further analyze the output of the system in order to look into the errors and phenomena underlying them. Shown below are some classes of errors that frequently occurred in our experiments.

We found that it is sometimes necessary to distinguish whether events happen remotely or not. In example (5), the system incorrectly links *at a Jewish temple* to both *avoid* and *disturbing* by the intra-sentential method. Event *disturbing* here occurs remotely with *worshippers* by context, and it must instead be linked to the location where *parade* occurs. However, it is very hard to find such fact on a surface level in this case.

- (5) The parade route was altered this year to [avoid EVENT] [disturbing EVENT] worshippers [at a Jewish temple LOC] during the observance of Yom Kippur.

Another important aspect is recognition of temporal relations between two events because temporal and spatial properties of events are highly correlated. Our system does not exploit temporal features for spatial relations, which is the limitation of the present work. In example (6), the system correctly recognizes *outside Iraq* for event *get* in the first sentence, but it passes on this link to event *getting* in the second sentence, erroneously linking it with *outside Iraq*, without considering the temporal distance between the two events. Note here that the two event words *get* and *getting* share a very high distributional similarity.

- (6) They were horrified by his injuries and pleaded for him to [get EVENT] specialist treatment [outside Iraq LOC]. His doctor, too, is grateful that Ali is [getting EVENT] this chance.

Most important, although it is obvious that our inter-sentential methods contribute to finding implicit event-location links and to improving overall performance, they still suffer from weakness in their behaviors: The information incorrectly identified in the previous steps is propagated to other neighboring events, producing a cascade of errors, especially when there is a significant temporal gap as shown above. It would be interesting to analyze additional factors for “a spatial transition” among multiple events. It would also be of help to exploit other contextual features of events, such as event coreference [Bejan and Harabagiu, 2014; Lu *et al.*, 2016], narrative schema [Chambers and Jurafsky, 2009], and sub-event hierarchy [Glavaš *et al.*, 2014] for detecting a valid sequence of spatial transitions.

6 Conclusion

In this paper, we addressed the problem of recognizing the implicit relations between events and locations on a document level, which has not been actively investigated compared to research on other event properties and relations. We started with the assumption that a series of events with high distributional similarities occur in a similar location. We showed that the distributional similarities are helpful for finding implicit event-location links especially when there are no explicit cues on a sentence level. Through the experimental analysis, we have shown that our method provides a good basis for this new task, despite the inherent difficulty of document-level inference. We believe that it would also be possible to ground spatial expressions identified by our system to a geolocation, which would be much more useful for location-based end applications.

Acknowledgements

This work was supported partly by Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (R0124-16-0002, Emotional Intelligence Technology to Infer Human Emotion and Carry on Dialogue Accordingly), and partly by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning (2017R1A2B4012788).

References

- [Barboza *et al.*, 2014] Philippe Barboza, Laetitia Vaillant, Yann Le Strat, David M. Hartley, Noele P. Nelson, Aba Mawudeku, Lawrence C. Madoff, Jens P. Linge, Nigel Collier, John S. Brownstein, and Pascal Astagneau. Factors influencing performance of internet-based biosurveillance systems used in epidemic intelligence for early detection of infectious diseases outbreaks. *PLoS ONE*, 9.3:e90536, 2014.
- [Bejan and Harabagiu, 2014] Cosmin Adrian Bejan and Sanda Harabagiu. Unsupervised event coreference resolution. *Computational Linguistics*, 40.2:311-347, 2014.
- [Chambers and Jurafsky, 2009] Nathanael Chambers and Dan Jurafsky. Unsupervised learning of narrative schemas and their participants. Joint Conference of the 47th Annual Meeting of the ACL and the 4th IJCNLP of the AFNLP, 2009.
- [Chung *et al.*, 2015] Jin-Woo Chung, Jinseon You, and Jong C. Park. Corpus annotation with a linguistic analysis of the associations between event mentions and spatial expressions. 29th Pacific Asia Conference on Language, Information and Computation, 2015.
- [D’Souza and Ng, 2015] Jennifer D’Souza and Vincent Ng. Sieve-based spatial relation extraction with expanding parse trees. *EMNLP 2015*.
- [Glavaš *et al.*, 2014] Goran Glavaš, Jan Šnajder, Marie-Francine Moens and Parisa Kordjamshidi. HiEve: A corpus for extracting event hierarchies from news stories. 9th LREC, 2014.
- [Hartley *et al.*, 2013] D. M. Hartley, N. P. Nelson, R. R. Arthur, P. Barboza, N. Collier, N. Lightfoot, J. P. Linge, E. van der Goot, A. Mawudeku, L. C. Madoff, L. Vaillant, R. Walters, R. Yangarber, J. Mantero, C. D. Corley, and J. S. Brownstein. An overview of internet biosurveillance. *Clinical Microbiology and Infection*, 19.11:1006–1013, 2013.
- [Hovy *et al.*, 2006] Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. OntoNotes: the 90% Solution. *HLT-NAACL 2006*.
- [Imran *et al.*, 2016] Muhammad Imran, Prasenjit Mitra, and Carlos Castillo. Twitter as a lifeline: Human-annotated twitter corpora for NLP of crisis-related messages. 10th LREC, 2016.
- [Kolomiyets *et al.*, 2013] Oleksandr Kolomiyets, Parisa Kordjamshidi, Steven Bethard, and Marie-Francine Moens. SemEval-2013 Task 3: Spatial role labeling, 7th International Workshop on Semantic Evaluation, 2013.
- [Kordjamshidi *et al.*, 2010] Parisa Kordjamshidi, Martijn van Otterlo, and Marie-Francine Moens. Spatial role labeling: Task definition and annotation scheme. 7th LREC, 2010.
- [Kordjamshidi *et al.*, 2012] Parisa Kordjamshidi, Steven Bethard, and Marie-Francine Moens. SemEval-2012 Task 3: Spatial role labeling. 6th International Workshop on Semantic Evaluation, 2012.
- [Lu *et al.*, 2016] Jing Lu, Deepak Venugopal, Vibhav Gogate, and Vincent Ng. Joint inference for event coreference resolution. 26th COLING, 2016.
- [Mani *et al.*, 2008] Inderjeet Mani, Janet Hitzeman, Justin Richer, Dave Harris, Rob Quimby, and Ben Wellner. SpatialML: Annotation scheme, corpora, and tools. 6th LREC, 2008.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. *NIPS 2013*.
- [Mírovský *et al.*, 2010] Jiří Mírovský, Lucie Mladová, and Šárka Zikánová. Connective-based measuring of the inter-annotator agreement in the annotation of discourse in PDT. 23rd COLING, 2010.
- [Mirza and Tonelli, 2016] Paramita Mirza, and Sara Tonelli. CATENA: CAusal and TEmporal relation extraction from NATural language texts. 26th COLING, 2016.
- [Pustejovsky *et al.*, 2003a] James Pustejovsky, José Castaño, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, and Graham Katz. TimeML: Robust specification of event and temporal expressions in text. 5th International Workshop on Computational Semantics, 2003.
- [Pustejovsky *et al.*, 2003b] James Pustejovsky, Patrick Hanks, Roser Saurí, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, and Marcia Lazo. The TIMEBANK corpus. *Corpus Linguistics 2003 Conference*.
- [Pustejovsky *et al.*, 2011] James Pustejovsky, Jessica Moszkowicz, and Marc Verhagen. ISO-Space: The annotation of spatial information in language. 6th Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation, 2011.
- [Pustejovsky *et al.*, 2015] James Pustejovsky, Parisa Kordjamshidi, Marie-Francine Moens, Aaron Levine, Seth Dworman, and Zachary Yocum. SemEval-2015 Task 8: SpaceEval. 9th International Workshop on Semantic Evaluation, 2015.
- [Roberts *et al.*, 2012] Kirk Roberts, Travis Goodwin, and Sanda Harabagiu. Annotating spatial containment relations between events. 8th LREC, 2012.
- [Roberts *et al.*, 2013] Kirk Roberts, Michael A. Skinner, and Sanda M. Harabagiu. Recognizing spatial containment relations between event mentions. 10th International Conference on Computational Semantics, 2013.
- [Vempala and Blanco, 2016] Alakananda Vempala and Eduardo Blanco. Beyond plain spatial knowledge: Determining where entities are and are not located, and for how long. 54th Annual Meeting of the ACL, 2016.