

A Large Scale Database of Strongly-related Events in Japanese

Tomohide Shibata, Shotaro Kohama and Sadao Kurohashi

Graduate School of Informatics, Kyoto University
Yoshida-honmachi, Sakyo-ku, Kyoto, 606-8501, Japan
{shibata, kohama, kuro}@nlp.ist.i.kyoto-u.ac.jp

Abstract

The knowledge about the relation between events is quite useful for coreference resolution, anaphora resolution, and several NLP applications such as dialogue system. This paper presents a large scale database of strongly-related events in Japanese, which has been acquired with our proposed method (Shibata and Kurohashi, 2011). In languages, where omitted arguments or zero anaphora are often utilized, such as Japanese, the coreference-based event extraction methods are hard to be applied, and so our method extracts strongly-related events in a two-phrase construct. This method first calculates the co-occurrence measure between predicate-arguments (events), and regards an event pair, whose mutual information is high, as strongly-related events. To calculate the co-occurrence measure efficiently, we adopt an association rule mining method. Then, we identify the remaining arguments by using case frames. The database contains approximately 100,000 unique events, with approximately 340,000 strongly-related event pairs, which is much larger than an existing automatically-constructed event database. We evaluated randomly-chosen 100 event pairs, and the accuracy was approximately 68%.

Keywords: event, case frames, knowledge acquisition

1. Introduction

Natural language understanding requires a wide variety of knowledge. One is the relation between predicate and argument. This relation has been automatically acquired in the form of case frames from a large corpus (Kawahara and Kurohashi, 2006; Materna, 2012), and is utilized for NLP tasks such as parsing. Another is the relation between events. The relation between events includes temporal relation, causality, and so on, and is useful for coreference resolution (Bean and Riloff, 2004), anaphora resolution (Gerber and Chai, 2010), and several NLP applications such as dialogue system.

Chambers and Jurafsky proposed a method for learning narrative schemas, which mean coherent sequences or sets of events, from unlabeled corpora (Chambers and Jurafsky, 2008; Chambers and Jurafsky, 2009). This method extracts two events that share a participant, called a *protagonist*. However, since this method relies on the coreference analysis result, they are hard to be applied to languages, such as Japanese, where omitted arguments or zero anaphora are often utilized.

Aiming to learn event knowledge in such languages, we previously proposed a method for automatically acquiring strongly-related events in Japanese (Shibata and Kurohashi, 2011). Our proposed method extracts strongly-related events in a two-phrase construct. First, since the arguments that specify the meaning of the predicate occur in at least one predicate-argument structure, the co-occurrence measure can be calculated from their occurrences. Thus, we can regard an event pair, whose mutual information is high, as strongly-related events. Next, we identify the remaining arguments by using case frames (Kawahara and Kurohashi, 2006).

This paper presents a large scale database of strongly-related events in Japanese, which has been acquired with the method (Shibata and Kurohashi, 2011). The database contains approximately 100,000 unique events, with approximately 340,000 strongly-related event pairs, which

is much larger than an existing automatically-constructed event database (Chambers and Jurafsky, 2010).

The rest of this paper is organized as follows: Section 2. reviews related work. Section 3. describes an overview of our proposed method (Shibata and Kurohashi, 2011). Section 4. describes our constructed resource.

2. Related Work

There are two types of resources of events: one is manually-constructed resource, and the other is a resource automatically constructed from a large corpus.

2.1. Manually-constructed Resource

Singh and Williams constructed a common sense knowledge base, called LifeNet, concerned with ordinary human activity (Singh and Williams, 2003). The knowledge base consists of 80,000 propositions with 415,000 temporal and atemporal links between propositions. Its scale is almost the same as our constructed resource. Espinosa and Lieberman proposed an EventNet, a toolkit for inferring temporal relations between commonsense events from the Openmind Commonsense Knowledge Base (Espinosa and Lieberman, 2005).

Recently, Regneri et al. collect natural language descriptions from volunteers over the Internet, and compute a temporal script graph (Regneri et al., 2010). They collected 493 event sequence descriptions for the 22 scenarios such as “eating in a fast-food restaurant” using the Amazon Mechanical Turk.

2.2. Automatically-constructed Resource

Chambers and Jurafsky learn narrative schemas, which mean coherent sequences or sets of events, from unlabeled corpora (Chambers and Jurafsky, 2008; Chambers and Jurafsky, 2009). This method extracts two events that share a participant, called a *protagonist*. They have made the constructed database publicly available (Chambers and Jurafsky, 2010). The database contains approximately 5,000 unique events combined into schemas of varying sizes.




Figure 1: An example of strongly-related events. (*ga* (nom), *wo* (acc), and *ni* (dat) are Japanese case markers.)

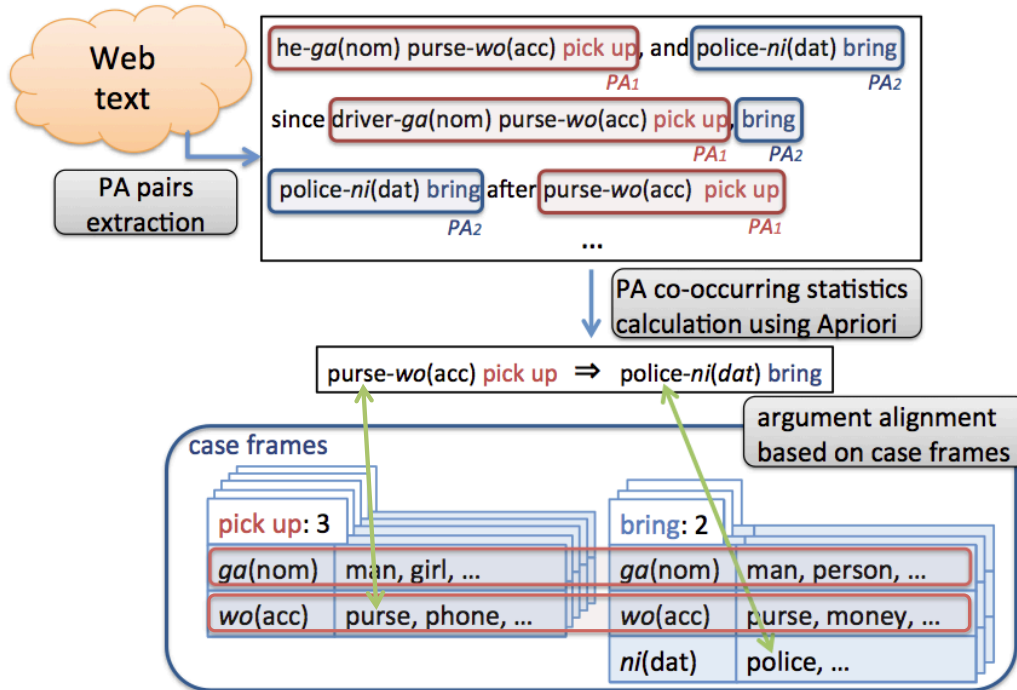


Figure 2: An overview of our proposed method.

Kasch and Oates proposed a method for extracting script-like structures from collections of Web documents (Kasch and Oates, 2010). Their method is topic-driven, and the experiment was performed on only one situation *eating at a restaurant*.

3. Strongly-Related Event Extraction

Our method focuses on Japanese, and extracts two strongly-related events in the form as shown in Figure 1. Figure 2 depicts an overview of our proposed method.

3.1. Predicate-Argument Structure Pairs Extraction

Strongly-related events appear in the form where they have a dependency relation with a variety of expressions (especially clause relation) in a text. For example, the event “*saifu*(purse)-*wo hirou*(pick up)” and the event “*keisatsu*(police)-*ni todokeru*(bring)” appear as follows:

- (1) *saifu-wo hiro-te keisatsu-ni todoke-ta*
 purse-acc pick up and police-dat brought
 ((A man) picked up a purse, and brought it to a police.)
- (2) *saifu-wo hiro-tta-node keisatsu-ni todoke-ta*
 purse-acc pick up because police-dat brought

class	words
77	<i>hachi</i> (bee), <i>ka</i> (mosquito), ...
105	dress, <i>ishou</i> (cloth), suit, ...
292	<i>keisatsu</i> (police), <i>kouan</i> (public peace), ...
502	address, <i>bangou</i> (number), ID, ...
752	<i>saifu</i> (purse), porch ...
956	<i>juugeki</i> (shooting), <i>shuugeki</i> (attack), ...
1829	<i>kenshuu</i> (training), intern, ...
1901	<i>douro</i> (road), <i>kokudou</i> (national highway), ...

Table 1: Examples of word class and its words.

(Because a man) picked up a purse, (he) brought it to a police.)

We extract two strongly-related events from a large number of pairs of two Predicate-Arguments (*PA*s) that have a dependency relation. From parsing results, a pair of *PA*s that have a dependency relation is first extracted. The extracted arguments are *ga* (nom), *wo* (acc), and *ni* (dat).

Argument Generalization

An argument is generalized to its word class so as to alleviate the problem of data sparseness. As a word class, a

large-scale clustering result of verb-noun dependency relations (Kazama and Torisawa, 2008) is used. The number of word class is 2,000, and this word class covers one million noun phrases. Table 1 shows examples of a word class and its words.

In pairs of the extracted PA_s s, the noun n is replaced with the word class $\langle c \rangle$ for which the probability $P(c|n)$ is maximal. For example, “ PA_1 : *saifu*(purse) *wo hirou*(pick up), PA_2 : *keisatsu*(police) *ni todokeru*(bring)” is changed to “ PA_1 : $\langle 752 \rangle$ *wo hirou*, PA_2 : $\langle 292 \rangle$ *ni todokeru*” since “*saifu*”, “*keisatsu*” belongs to the word class $\langle 752 \rangle$, $\langle 292 \rangle$, respectively. In the same way, “ PA_1 : porch *wo hirou*(pick up), PA_2 : *keisatsu*(police) *ni todokeru*(bring)” is changed to “ PA_1 : $\langle 752 \rangle$ *wo hirou*, PA_2 : $\langle 292 \rangle$ *ni todokeru*”, and thus, these two PA_s can be identical, which alleviates the problem of data sparseness.

3.2. Co-occurrence Statistics Calculation between Predicate-Argument Structures

Given a lot of PA_s s, as extracted in Section 3.1., the co-occurrence statistics between PA_s s is calculated. Since the number of pairs of arbitrary PA_s s is enormous, a question that arises is how to obtain related PA_s s effectively. To solve this problem, we adopt an *association rule mining* method (Agrawal et al., 1993) for the calculation of co-occurrence statistics between PA_s s. The association rule mining method can efficiently seek candidate items that satisfy specific conditions.

3.2.1. Association Rule Mining

Association rule mining is a method for discovering significant rules in a large database (Agrawal et al., 1993). This method is originally designed to discover rules such as “a customer who buys diapers tends to buy beer” in customer transactions.

Let $I = I_1, I_2, \dots, I_m$ be a set of binary attributes, called items. Transaction t is defined as a set of items ($t \subseteq I$), and transaction database T is defined as a set of transactions ($T = t_1, t_2, \dots, t_n$).

A rule is defined as an implication of the form $X \Rightarrow Y$ where $X, Y \subseteq I$ and $X \cap Y = \phi$. This signifies “if X occurs, Y tends to occur”. The set of items X and Y are called antecedent (left-hand side, lhs) and consequent (right-hand side, rhs) of the rule respectively. For every rule, the following three measures are defined:

$$\text{support}(X \Rightarrow Y) = \frac{C(X \cup Y)}{|T|} \quad (1)$$

$$\text{confidence}(X \Rightarrow Y) = \frac{\text{support}(X \Rightarrow Y)}{\text{support}(X)} \quad (2)$$

$$\text{lift}(X \Rightarrow Y) = \frac{\text{confidence}(X \Rightarrow Y)}{\text{support}(Y)}, \quad (3)$$

where $C(X)$ represents the number of transactions containing the item X .

The *support* is defined as the fraction formed the number of transactions that contain the itemset X and the total number of transactions in the database. The *confidence* is defined as the fraction formed from the transactions that contain $X \cup Y$ and the transactions that contain X . The *lift* corresponds to pointwise mutual information between X and Y .

Apriori algorithm (Borgelt and Kruse, 2002) is one of the well-known implementations for association rule mining. This algorithm exploits the observation that no superset of an infrequent itemset can be frequent, and uses breadth-first search and a tree structure to seek candidate items.

The input for Apriori algorithm is transaction data, the minimum support, and minimum confidence, and the algorithm enumerates all rules that satisfy the specified conditions.

3.2.2. Apriori Algorithm Application to Co-occurrence Calculation

The Apriori algorithm is applied to the calculation of co-occurrence statistics between PA_s s. An item introduced in Section 3.2.1. corresponds to a predicate or an argument, and a transaction is obtained from a pair of PA_s s. Examples of transaction data are shown in Table 2.

Since the rules we want to extract are supposed to satisfy the following conditions:

- X (left-hand side) consists of a predicate of PA_1 , and zero or more arguments in PA_1
- Y (right-hand side) consists of a predicate of PA_2 , and zero or more arguments in PA_2 ,

all the rules that do not satisfy these conditions are discarded. Among those that do, the rule for which the lift is higher than *lift-min* and less than *lift-max* is adopted. It is well-known that the pointwise mutual information (which corresponds to lift) for which the frequency is low gets extremely high, and thus rules for which the lift is greater than *lift-max* are discarded.

The Apriori algorithm naturally judges which argument is relevant for each predicate pair. For example, from the transaction data shown in Table 2, the following rule is obtained:

1. *saifu-wo hirou* \Rightarrow *keisatsu-ni todokeru*
2. *saifu-wo hirou* \Rightarrow *tewatasu*

The first rule implies that for the predicate pair “*hirou*” and “*todokeru*”, “*saifu-wo*” for the predicate in PA_1 and “*keisatsu-ni*” for the predicate in PA_2 are relevant. Similarly, the second rule implies that for the predicate pair “*hirou*” and “*tewatasu*”, “*saifu-wo*” for the predicate in PA_1 is relevant.

3.3. Argument Alignment based on Case Frames

As mentioned in Introduction, since an argument is often omitted in the extracted predicate-argument pairs, there is usually a lack of arguments in the extracted rules as described in the previous section. In the following rule, the argument of the *wo* case in PA_1 corresponds to the *wo* case in PA_2 , and the argument that includes nouns such as “*otoko*(man)”, “*hito*(person)” acts for the *ga* case both in PA_1 and PA_2 .

$$\text{saifu-wo hirou} \Rightarrow \text{keisatsu-ni todokeru}$$

Such alignment between arguments can be performed by case frames. The case frames are constructed automatically by clustering similar predicate usages from a raw corpus, and thus each predicate has several case frames. Examples

PA_1		PA_2	
arguments	predicate	arguments	predicate
<i>saifu</i> (purse)- <i>wo</i>	<i>hirou</i> (pick up)	<i>keisatsu</i> (police)- <i>ni</i>	<i>todokeru</i> (bring)
<i>kare</i> (he)- <i>ga</i> , <i>saifu</i> - <i>wo</i>	<i>hirou</i>	<i>keisatsu</i> - <i>ni</i>	<i>todokeru</i>
<i>saifu</i> - <i>wo</i>	<i>hirou</i>		<i>todokeru</i>
	<i>hirou</i>	<i>keisatsu</i> - <i>ni</i>	<i>todokeru</i>
		...	
<i>saifu</i> - <i>wo</i>	<i>hirou</i>		<i>tewatasu</i> (hand)
<i>saifu</i> - <i>wo</i>	<i>hirou</i>	<i>kare</i> (he)- <i>ni</i>	<i>tewatasu</i>
<i>otoko</i> (man)- <i>ga</i> , <i>saifu</i> - <i>wo</i>	<i>hirou</i>		<i>tewatasu</i>
		...	

Table 2: Examples of transaction data. (One line represents a transaction.)

verb	case marker	examples
<i>hirou</i> :1 (pick up)	<i>ga</i> <i>wo</i>	<i>josei</i> (lady), <i>hito</i> (person), ... taxi, <i>kuruma</i> (car), ...
		...
<i>hirou</i> :3 (pick up)	<i>ga</i> <i>wo</i>	<i>otoko</i> (man), <i>onnanoko</i> (girl), ... <i>saifu</i> (purse), <i>denwa</i> (phone) ...
		...
<i>todokeru</i> :1 (deliver)	<i>ga</i> <i>wo</i>	staff, <i>syokuin</i> (staff), ... <i>kyohou</i> (information), news, ...
		...
<i>todokeru</i> :2 (bring)	<i>ga</i> <i>wo</i> <i>ni</i>	<i>otoko</i> (man), <i>hito</i> (person), ... <i>saifu</i> (purse), <i>kane</i> (money), ... <i>keisatsu</i> (police), ...
		...

Table 3: Examples of the automatically constructed case frames.

of the case frames are shown in Table 3. When both a case in cf_1 assigned to PA_1 and a case in cf_2 assigned to PA_2 have a similar distribution of examples, the case in PA_1 and the case in PA_2 can be aligned.

The best combinations of the case frame in both PA_1 and PA_2 and the best alignment of cases are determined as follows:

1. If there is an argument, select case frames corresponding to the argument, otherwise, all case frames are candidates. In the above example, while in PA_1 the case frame 3 is selected according to the argument for the case *wo* (“*saifu*”), in PA_2 the case frame 2 is selected according to the case *ni* (“*keisatsu*”).
2. Choose the best case frame pairs that maximize the following score:

$$\operatorname{argmax}_{cf_1, cf_2} \max_{\mathbf{a}} \sum_{a \in \mathbf{a}} \operatorname{sim}(arg_1, a(arg_1)) \quad (4)$$

where \mathbf{a} denotes the alignment of case components between PA_1 and PA_2 , arg_1 denotes an argument in PA_1 , $a(arg_1)$ denotes an argument in PA_2 that aligned with arg_1 , and sim denotes the cosine similarity of the case components distribution between arg_1

extracted rule	correct		incorrect
	85(85.0%)		15(15.0%)
argument alignment	correct	incorrect	
	58(68.2%)	27(31.8%)	–

Table 4: Accuracy of extracted rule and the argument alignment.

and $a(arg_1)$. In the example, the alignment between the case *ga* of the case frame 3 in PA_1 and the case *ga* of the case frame 2 in PA_2 , and the case *wo* in PA_1 and the case *wo* in PA_2 is performed.

4. Constructed Resource

Approximately 100 million Japanese Web pages were used to extract strongly-related events. These pages include 6 billion sentences, containing 100 billion words. Owing to the presence of many duplicate pages on the Web, such as mirror pages, duplicate sentences were discarded. Thus, 1.6 billion sentences containing approximately 25 billion words were acquired. The case frames were automatically constructed from the same Web corpus with a method proposed by (Kawahara and Kurohashi, 2006).

The Web corpus was processed using the Japanese Morphological Analyzer JUMAN¹ and the Japanese parser KNP², and pairs of PA_s were extracted.

In the application of Apriori algorithm presented in Section 3.2., the minimum support, confidence was set to 3.0×10^{-8} , 3.0×10^{-4} respectively, and *lift-min*, *lift-max* was set to 10, 10,000 respectively.

The acquired database contains approximately 100,000 unique events, with approximately 340,000 strongly-related event pairs, which is much larger than an existing automatically-constructed event database (Chambers and Jurafsky, 2010).

4.1. Evaluation

We chose 100 rules at random, and evaluated whether each is valid. The upper part in Table 4 shows the accuracy, and we found 85 valid rules of the 100 rules, and the accuracy was 0.85.

¹<http://nlp.ist.i.kyoto-u.ac.jp/EN/index.php?JUMAN>

²<http://nlp.ist.i.kyoto-u.ac.jp/EN/index.php?KNP>

	PA_1		PA_2		evaluation
	argument	predicate	argument	predicate	
(1)	$A_1: \{ \textit{boshuu}, \textit{moushikomi}, \dots \} \textit{ga}$ (invitation) (application)	$\textit{tassuru}$ (reach)	$\Rightarrow A_1: \{ \textit{boshuu}, \textit{moushikomi}, \dots \} \textit{wo}$ (invitation) (application)	$\textit{shimekiru}$ (close)	correct
	$A_2: \{ \textit{teiin} \}$ (capacity)				
(2)	$A_1: \{ \textit{watashi}, \textit{kodomo}, \dots \} \textit{ga}$ (I) (child)	$\textit{sotsugyo}$ (graduate)	$\Rightarrow A_1: \{ \textit{watashi}, \textit{kodomo}, \dots \} \textit{ga}$ (I) (child)	$\textit{shuusyoku}$ (get a job)	correct
	$A_2: \{ \textit{daigaku} \}$ (university)		$A_3: \{ \textit{kaisha} \}$ (company)		
(3)	$A_1: \{ \textit{musuko}, \textit{kodomo}, \dots \} \textit{ga}$ (son) (child)	\textit{tentou} (fall down)	$\Rightarrow A_1: \{ \textit{musuko}, \textit{kodomo}, \dots \} \textit{ga}$ (son) (child)	$\textit{kossetsu}$ (fracture)	correct
(4)	$A_1: \{ \textit{sakuhin}, \dots \} \textit{ga}$ (product)	$\textit{nominatesareru}$ (nominate)	$\Rightarrow A_1: \{ \textit{sakuhin}, \dots \} \textit{ga}$ (product)	\textit{jusyuo} (win an award)	correct
	$A_2: \{ \textit{sho}, \textit{yuushuu-sho}, \dots \} \textit{ni}$ (prize) (grand prix)		$A_2: \{ \textit{sho}, \textit{yuushuu-sho}, \dots \} \textit{wo}$ (prize) (grand prix)		
(5)	$A_1: \{ \textit{watashi}, \textit{hito}, \dots \} \textit{ga}$ (I) (person)	$\textit{tazuneru}$ (visit)	$\Rightarrow A_1: \{ \textit{watashi}, \textit{hito}, \dots \} \textit{ga}$ (I) (person)	\textit{ukagau} (hear)	correct
	$A_2: \{ \textit{sensei}, \textit{shachou}, \dots \} \textit{wo}$ (teacher) (chief)		$\Rightarrow A_2: \{ \textit{sensei}, \textit{shachou}, \dots \} \textit{ni}$ (teacher) (chief)		
			$A_3: \{ \textit{hanashi}, \dots \} \textit{wo}$ (talk)		
(6)	$A_1: \{ \textit{kanojo}, \textit{josei}, \dots \} \textit{ga}$ (she) (lady)	$\textit{purezento}$ (present)	$\Rightarrow A_2: \{ \textit{shouhin}, \textit{hana}, \dots \} \textit{ga}$ (goods) (flower)	$\textit{yorokobareru}$ (delighted)	incorrect
	$A_2: \{ \textit{shouhin}, \textit{hana}, \dots \} \textit{wo}$ (goods) (flower)		$A_1: \{ \textit{kanojo}, \textit{josei}, \dots \} \textit{ni}$ (she) (lady)		
(7)	$A_1: \{ \textit{kodomo}, \dots \} \textit{ga}$ (child)	\textit{kekkon} (get married)	$\Rightarrow A_1: \{ \textit{kodomo}, \dots \} \textit{ga}$ (child)	\textit{iru} (have)	incorrect

Table 5: Examples of acquired strongly-related events. (The underlined arguments indicate the one acquired by the association rule mining method.)




Figure 3: The correct alignment of (6) in Table 5.

Then, We calculated the accuracy of the argument alignment for the 85 rules. The bottom part in Table 4 shows the accuracy, and we found 58 of 85 were valid, and the accuracy was 0.682. Table 5 shows examples of acquired strongly-related events.

A major error is a case alignment error where the case component distribution between two cases in a PA is very similar. In the example (6), the alignment shown in Figure 3 is correct. This error was caused by the fact that the case ga and the case ni in PA_1 and the case ga and the case ni in PA_2 include nouns representing an agent.

We are going to make the acquired event knowledge publicly available. Since the accuracy of the constructed resources may not be sufficient to be utilized for other tasks or applications, we are also going to modify wrong event pairs using crowdSourcing, as Zeichner et al. (Zeichner et al., 2012).

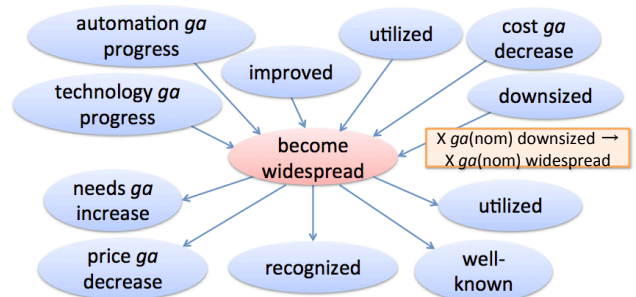


Figure 4: Network structure between events concerned with “become widespread”.

4.2. Event Network Structure

Figure 4 is an example of a network structure between events concerned with “become widespread”, which is constructed from strongly-related events obtained by our pro-

posed method. As we can see, several events were acquired before/after the event “become widespread” occurs.

5. Conclusion

This paper presented a large scale database of strongly-related events in Japanese, which has been acquired with the method (Shibata and Kurohashi, 2011). This method first calculated the co-occurrence measure between predicate-arguments (events), and regarded an event pair, whose mutual information is high, as strongly-related events. To calculate the co-occurrence measure efficiently, we adopted association rule mining. Then, we identified the remaining arguments by using case frames. The database contains approximately 100,000 unique events, with approximately 340,000 strongly-related event pairs, which is much larger than an existing automatically-constructed event database.

We are going to improve the accuracy of the automatic extraction method as well as modify wrong acquired event pairs using crowdSourcing. Furthermore, we are going to demonstrate the usefulness of the acquired event pairs in the application of anaphora resolution.

6. References

- Agrawal, R., Imielinski, T., and Swami, A. (1993). Mining association rules between sets of items in large databases. In *Proceedings of the ACM-SIGMOD 1993 International Conference on Management of Data (1993)*, pages 207–216.
- Bean, D. and Riloff, E. (2004). Unsupervised learning of contextual role knowledge for coreference resolution. In *HLT-NAACL 2004: Main Proceedings*, pages 297–304.
- Borgelt, C. and Kruse, R. (2002). Induction of association rules: Apriori implementation. In *Proceedings of 15th Conference on Computational Statistics*, pages 395–400.
- Chambers, N. and Jurafsky, D. (2008). Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797.
- Chambers, N. and Jurafsky, D. (2009). Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610.
- Chambers, N. and Jurafsky, D. (2010). A database of narrative schemas. In *LREC-2010*.
- Espinosa, J. and Lieberman, H. (2005). EventNet: Inferring temporal relations between commonsense events. In *Proceedings of the 4th Mexican International Conference on Artificial Intelligence*, pages 61–69.
- Gerber, M. and Chai, J. (2010). Beyond NomBank: A study of implicit arguments for nominal predicates. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1583–1592.
- Kasch, N. and Oates, T. (2010). Mining script-like structures from the web. In *Proceedings of the NAACL HLT 2010 First International Workshop on Formalisms and Methodology for Learning by Reading*, pages 34–42.
- Kawahara, D. and Kurohashi, S. (2006). A fully-lexicalized probabilistic model for Japanese syntactic and case structure analysis. In *Proceedings of the HLT-NAACL2006*, pages 176–183.
- Kazama, J. and Torisawa, K. (2008). Inducing gazetteers for named entity recognition by large-scale clustering of dependency relations. In *Proceedings of ACL-08: HLT*, pages 407–415.
- Materna, J. (2012). LDA-Frames: An Unsupervised Approach to Generating Semantic Frames. In *Proceedings of the 13th International Conference CICLing 2012, Part I*, pages 376–387.
- Regneri, M., Koller, A., and Pinkal, M. (2010). Learning script knowledge with web experiments. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 979–988.
- Shibata, T. and Kurohashi, S. (2011). Acquiring strongly-related events using predicate-argument co-occurring statistics and case frames. In *Proceedings of the 5th International Joint Conference on Natural Language Processing (IJCNLP2011, poster)*, pages 1028–1036.
- Singh, P. and Williams, W. (2003). Lifenet: A propositional model of ordinary human activity. In *Proceedings of Workshop on Distributed and Collaborative Knowledge Capture*.
- Zeichner, N., Berant, J., and Dagan, I. (2012). Crowdsourcing inference-rule evaluation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012)*, pages 156–160.