Acquiring Event Relation Knowledge by Learning Cooccurrence Patterns and Fertilizing Cooccurrence Samples with Verbal Nouns

Shuya Abe **Kentaro Inui** **Yuji Matsumoto**
Graduate School of Information Science,
Nara Institute of Science and Technology
{shuya-a,inui,matsu}@is.naist.jp

Abstract

Aiming at acquiring semantic relations between events from a large corpus, this paper proposes several extensions to a state-of-the-art method originally designed for entity relation extraction, reporting on the present results of our experiments on a Japanese Web corpus. The results show that (a) there are indeed specific cooccurrence patterns useful for event relation acquisition, (b) the use of cooccurrence samples involving verbal nouns has positive impacts on both recall and precision, and (c) over five thousand relation instances are acquired from a 500M-sentence Web corpus with a precision of about 66% for *action-effect* relations.

1 Introduction

The growing interest in practical NLP applications such as question answering, information extraction and multi-document summarization places increasing demands on the processing of relations between textual fragments such as entailment and causal relations. Such applications often need to rely on a large amount of lexical semantic knowledge. For example, a causal (and entailment) relation holds between the verb phrases *wash something* and *something is clean*, which reflects the commonsense notion that if someone has washed something, this object is clean as a result of the washing event. A crucial issue is how to obtain and maintain a potentially huge collection of such event relations instances.

Motivated by this background, several research groups have reported their experiments on automatic acquisition of causal, temporal and entailment relations between event mentions (typically verbs or verb phrases) (Lin and Pantel, 2001; Inui et al., 2003; Chklovski and Pantel, 2005; Torisawa, 2006; Pekar, 2006; Zanzotto et al., 2006, etc.). The common idea behind them is to use a small number of manually selected generic lexico-syntactic cooccurrence patterns (LSPs or simply patterns). *to Verb-X and then Verb-Y*, for example, is used to obtain temporal relations such as *marry* and *divorce* (Chklovski and Pantel, 2005). The use of such generic patterns, however, tends to be high recall but low precision, which requires an additional component for pruning extracted relations. This issue has been addressed in basically two approaches, either by devising heuristic statistical scores (Chklovski and Pantel, 2005; Torisawa, 2006; Zanzotto et al., 2006) or training classifiers for disambiguation with heavy supervision (Inui et al., 2003).

This paper explores a third way for enhancing present LSP-based methods for event relation acquisition. The basic idea is inspired by the following recent findings in relation extraction (Ravichandran and Hovy, 2002; Pantel and Pennacchiotti, 2006, etc.), which aims at extracting semantic relations between *entities* (as opposed to *events*) from texts. (a) The use of generic patterns tends to be high recall but low precision, which requires an additional component for pruning. (b) On the other hand, there are specific patterns that are highly reliable but they are much less frequent than generic patterns and each makes only a small contribution to recall. (c) Combining a few generic patters with a much larger collection of reliable specific patterns boosts both precision and recall. Such specific patterns can be acquired from a very large corpus with seeds.

Given these insights, an intriguing question is whether the same story applies to event relation acquisition as well or not. In this paper, we explore this issue through the following steps. First, while previous methods use only verb-verb cooccurrences, we use cooccurrences between verbal nouns and verbs such as *cannot* ⟨*find out (something)*⟩ *due to the lack of* ⟨*investigation*⟩ as well as verb-verb cooccurrences. This extension dramatically enlarge the pool of potential candidate LSPs (Section 3). Second, we extend Pantel and Pennacchiotti (2006)'s Espresso algorithm, which induces specific reliable LSPs in a bootstrapping manner for entity relation extraction, so that the extended algorithm can apply to event relations (Sections 4). Third, we report on the present results of our empirical experiments, where the extended algorithm is applied to a Japanese 500M-sentence Web corpus to acquire two types of event relations, *action-effect* and *action-means* relations (Section 5)

2 Related work

Perhaps a simplest way of using LSPs for event relation acquisition can be seen in the method Chklovski and Pantel (2005) employ to develop VerbOcean. Their method uses a small number of manually selected generic LSPs such as *to Verb-X and then Verb-Y* to obtain six types of semantic relations including *strength* (e.g. *taint – poison*) and *happens-before* (e.g. *marry – divorce*) and obtain about 29,000 verb pairs with 65.5% precision.

One way for pruning extracted relations is to incorporate a classifier trained with supervision. Inui et al. (2003), for example, use a Japanese generic causal connective marker *tame* (because) and a supervised classifier learner to separately obtain four types of causal relations: *cause*, *precondition*, *effect* and *means*.

Torisawa (2006), on the other hand, acquires entailment relations by combining the verb pairs extracted with a highly generic connective pattern *Verb-X and Verb-Y* together with the cooccurrence statistics between verbs and their arguments. While the results Torisawa reports look promising, it is not clear yet if the method applies to other types of relations because it relies on relation-specific heuristics.

Another direction from (Chklovski and Pantel, 2005) is in the use of LSPs involving nominalized verbs. Zanzotto et al. (2006) obtain, for example, an entailment relation *X wins* → *X plays* from such a pattern as *player wins*. However, their way of using nominalized verbs is highly limited compared with our way of using verbal nouns.

3 Espresso

This section overviews Pantel and Pennacchiotti (2006)'s Espresso algorithm. Espresso takes as input a small number of seed instances of a given target relation and iteratively learns cooccurrence patterns and relation instances in a bootstrapping manner.

Ranking cooccurrence patterns For each given relation instance $\{x, y\}$, Espresso retrieves the sentences including both x and y from a corpus and extracts from them cooccurrence samples. For example, given an instance of the *is-a* relation such as ⟨*Italy,country*⟩, Espresso may find cooccurrence samples such as *countries such as Italy* and extract such a pattern as *Y such as X*. Espresso defines the reliability $r_\pi(p)$ of pattern p as the average strength of its association with each relation instance i in the current instance set I, where each instance i is weighted by its reliability $r_\iota(i)$:

$$r_\pi(p) = \frac{1}{|I|} \sum_{i \in I} \frac{pmi(i,p)}{max_{pmi}} \times r_\iota(i) \qquad (1)$$

where $pmi(i,p)$ is the pointwise mutual information between i and p, and max_{pmi} is the maximum PMI between all patterns and all instances.

Ranking relation instances Intuitively, a reliable relation instance is one that is highly associated with multiple reliable patterns. Hence, analogously to the above pattern reliability measure, Espresso defines the reliability $r_\iota(i)$ of instance i as:

$$r_\iota(i) = \frac{1}{|P|} \sum_{p \in P} \frac{pmi(i,p)}{max_{pmi}} \times r_\pi(p) \qquad (2)$$

where $r_\pi(p)$ is the reliability of pattern p, defined above in (1), and max_{pmi} is as before. $r_\iota(i)$ and $r_\pi(p)$ are recursively defined, where $r_\iota(i) = 1$ for each manually supplied seed instance i.

4 Event relation acquisition

Our primary concerns are whether there are indeed specific cooccurrence patterns useful for acquiring event relations and whether such patterns can be found in a bootstrapping manner analogous to Espresso. To address these issues, we make several extensions to Espresso, which is originally designed for entity relations (not scoping event relations).

4.1 Cooccurences with verbal nouns

Most previous methods for event relation acquisition rely on verb-verb cooccurrences because verbs (or verb phrases) are the most typical device for referring to events. However, languages have another large class of words for event reference, namely verbal nouns or nominalized forms of verbs. In Japanese, for example, verbal nouns such as *kenkyu* (research) constitute the largest morphological category used for event reference.

Japanese verbal nouns have dual statuses, as verbs and nouns. When occurring with the verb *suru* (do-NON-PAST), verbal nouns function as a verb as in (1a). On the other hand, when accompanied by case markers such as *ga* (NOMINATIVE) and *o* (ACCUSATIVE), they function as a noun as in (1b). Finally, but even more importantly, when accompanied by a large variety of suffixes, verbal nouns constitute compound nouns highly productively as in (1c).

(1) a. *Ken-ga* *gengo-o* *kenkyu-suru*
Ken-NOM language-ACC research-PRES
Ken researches on language.

b. *Ken-ga* *gengo-no* *kenkyu-o* *yame-ta*
Ken-NOM language-on research-ACC quit-PAST
Ken quitted research on language.

c. *-sha* (person):
e.g. *kenkyu-sha* (researcher)
-shitsu (place):
e.g. *kenkyu-shitsu* (laboratory)
-go (after):
e.g. *kenkyu-go* (after research)

These characteristics of verbal nouns can be made use of to substantially increase both cooccurrence instances and candidate cooccurrence patterns (see Section 5.1 for statistics). For example, the verbal noun *kenkyu* (research) often cooccurs with the verb *jikken* (experiment) in the pattern of (2a). From those cooccurrences, one may learn that *jikken-suru* (to experiment) is an action that is often taken as a part of *kenkyu-suru* (to research). In such a case, we may consider a pattern as shown in (2b) useful for acquiring *part-of* relations between actions.

(2) a. *kenkyu-shitsu-de* *jikken-suru*
research-place-in experiment-VERB
conduct experiments in the laboratory

b. *(Act-X)-shitsu-de* *(Act-Y)-suru*
(Act-X)-place-in (Act-X)-VERB
(Act-Y) is often done in doing (Act-X)

When functioning as a noun, verbal nouns are potentially ambiguous between the event reading and the entity/object reading. For example, the verbal noun *denwa* (phone) in the context *denwa-de* (phone-by) may refer to either a phone-call event or a physical phone. While, ideally, such eventhood ambiguities should be resolved before collecting cooccurrence samples with verbal nouns, we simply use all the occurrences of verbal nouns in collecting cooccurrences in our experiments. It is an interesting issue for future work whether eventhood determination would have a strong impact on the performance of event relation extraction.

4.2 Selection of arguments

One major step from the extraction of entity relations to the extraction of event relations is how to address the issue of *generalization*. In entity relation extraction, relations are typically assumed to hold between chunks like named entities or simply between one-word terms, where the issue of determining the appropriate level of the generality of extracted relations has not been salient. In event relation extraction, on the other hand, this issue immediately arises. For example, the cooccurrence sample in (3) suggests the *action-effect* relation between *niku-o yaku* (grill the meat) and *(niku-ni) kogeme-ga tsuku* ((the meat) gets brown)[1].

(3) *(kogeme-ga* *tsuku) -kurai* *niku-o* *yaku*
a burn-NOM get -so that meat-ACC grill
grill the meat so that it gets brown
(grill the meat to a deep brown)

[1] The parenthesis in the first row of (3) indicates a subordinate clause.

In this relation, the argument *niku* (meat) of the verb *yaku* (grill) can be dropped and generalized to *something to grill*; namely the *action-effect* relation still holds between *X-o yaku* (grill X) and *X-ni kogeme-ga tsuku* (X gets brown). On the other hand, however, the argument *kogeme* (a burn) of the verb *tsuku* (get) cannot be dropped; otherwise, the relation would no longer hold.

One straightforward way to address this problem is to expand each cooccurrence sample to those corresponding to different degrees of generalization and feed them to the relation extraction model so that its scoring function can select appropriate event pairs from expanded samples. For example, cooccurrence sample (3) is expanded to those as in (4):

(4) a. *(kogeme-ga* *tsuku) -kurai* *niku-o* *yaku*
a burn-NOM get -so that meat-ACC grill

b. *(tsuku) -kurai* *niku-o* *yaku*
get -so that meat-ACC grill

c. *(kogeme-ga* *tsuku) -kurai* *yaku*
a burn-NOM get -so that grill

d. *(tsuku) -kurai* *yaku*
get -so that grill

In practice, in our experiments (Section 5), we restrict the number of arguments for each event up to one to avoid the explosion of the types of infrequent candidate relation instances.

4.3 Volitionality of events

Inui et al. (2003) discuss how causal relations between events should be typologized for the purpose of semantic inference and classify causal relations basically into four types — Effect, Means, Precondition and Cause relations — based primarily on the volitionality of involved events. For example, Effect relations hold between volitional actions and their resultative non-volitional states/happenings/experiences, while Cause relations hold between only non-volitional states/happenings/experiences.

Following this typology, we are concerned with the volitionality of each event mention. For our experiments, we manually built a lexicon of over 12,000 verbs (including verbal nouns) with volitionality labels, obtaining 8,968 volitional verbs, 3,597 non-volitional and 547 ambiguous. Volitional verbs include *taberu* (eat) and *kenkyu-suru* (research), while non-volitional verbs include *atatamaru* (get warm), *kowareru* (to break-*vi*) and *kanashimu* (be sad). We discarded the ambiguous verbs in the experiments.

4.4 Dependency-based cooccurrence patterns

The original Espresso encodes patterns simply as a word sequence because entity mentions in the relations it scopes tend to cooccur locally in a single phrase or clause. In event relation extraction, however, cooccurrence patterns of event mentions in the relations we consider (causal relations, temporal relations, etc.) can be captured better as a path on a syntactic dependency tree because (i) such mention pairs tend to cooccur in a longer dependency path and (ii) as discussed in Section 4.2, we want to exclude the arguments of event mentions from cooccurrence patterns, which would be difficult with word sequence-based representations of patterns.

A Japanese sentence can be analyzed as a sequence of base phrase (BP) chunks called *bunsetsu* chunks, each which typically consists of one content (multi-)word followed by functional words. We assume each sentence of our corpus is given a dependency parse tree over its BP chunks. Let us call a BP chunk containing a verb or verbal noun an *event chunk*. We create a cooccurrence sample from any pair of event chunks that cooccur if either (a) one event chunk depends directly on the other, or (b) one event chunk depends indirectly on the other via one intermediate chunk. Additionally, we apply the Japanese functional expressions dictionary(Matsuyoshi et al., 2006) to a cooccurrence pattern for generalization.

In (5), for example, the two event chunks, *taishoku-go-ni* (after retirement) and *hajimeru* (begin), meet the condition (b) above and the dependency path designated by bold font is identified as a candidate cooccurrence pattern. The argument *PC-o* of the verb *hajimeru* is excluded from the path.

(5) *(taishoku-**go-no*** ***tanoshimi)-ni*** *PC-o* *hajimeru*
retirement-after as a hobby PC-ACC begin
begin a PC as a hobby after retirement

Table 1: Examples of acuired cooccurrence patterns and relatio instances for the action-effect relation

freq	cooccurrence patterns	relation instances
94477	$\langle$*verb;action*$\rangle$***temo***$\langle$*verb;effect*$\rangle$**nai** (to do $\langle$action$\rangle$ though $\langle$effect$\rangle$ dose not happen)	*sagasu::mitsukaru* (search::be found), *asaru::mitsukaru* (hunt::be found), *purei-suru::kuria-suru* (play::finish)
6250	$\langle$*verb;action*$\rangle$***takeredomo***$\langle$*verb;effect*$\rangle$**nai** (to do $\langle$action$\rangle$ though $\langle$effect$\rangle$ dose not happen)	*shashin-wo-toru::toreru* (shot photograph::be shot), *meiru-wo-okuru::henji-ga-kaeru* (send a mail::get an answer)
1851	$\langle$*noun;action*$\rangle$***wo-shitemo***$\langle$*verb;effect*$\rangle$**nai** (to do $\langle$action$\rangle$ though $\langle$effect$\rangle$ dose not happen)	*setsumei-suru::nattoku-suru* (explain::agree), *siai-suru::katsu* (play::win), *siai-suru::makeru* (play::lose)
1329	$\langle$*verb;action*$\rangle$***yasukute***$\langle$*adjective;effect*$\rangle$ (to simply do $\langle$action$\rangle$ and $\langle$effect$\rangle$)	*utau::kimochiyoi* (sing::feel good), *hashiru::kimochiyoi* (run::feel good)
4429	$\langle$*noun;action*$\rangle$***wo-kiite***$\langle$*verb;effect*$\rangle$ (to hear $\langle$action$\rangle$ so that $\langle$effect$\rangle$)	*setsumei-suru::nattoku-suru* (explain::agree), *setsumei-suru::rikai-dekiru* (explain::can understand)

5 Experiments

5.1 Settings

For an empirical evaluation, we used a sample of approximately 500M sentences taken from the Web corpus collected by Kawahara and Kurohashi (2006). The sentences were dependency-parsed with Cabocha (Kudo and Matsumoto, 2002), and cooccurrence samples of event mentions were extracted. Event mentions with patterns whose frequency was less than 20 were discarded in order to reduce computational costs. As a result, we obtained 34M cooccurrence tokens with 11M types. Note that among those cooccurrence samples 15M tokens (44%) with 4.8M types (43%) are those with verbal nouns, suggesting the potential impacts of using verbal nouns.

In our experiments, we considered two of Inui et al. (2003)'s four types of causal relations: *action-effect* relations (Effect in Inui et al.'s terminology) and *action-means* relations (Means). An *action-effect* relation holds between events x and y if and only if non-volitional event y is likely to happen as either a direct or indirect effect of volitional action x. For example, the action *X-ga undou-suru* (X exercises) and the event *X-ga ase-o kaku* (X sweats) are considered to be in this type of relation. A *action-means* relation holds between events x and y if and only if volitional action y is likely to be done as a part/means of volitional action x. For example, if case a event-pair is *X-ga hashiru* (X runs) is considered as a typical action that is often done as a part of the action *X-ga undou-suru* (X exercises).

Note that in these experiments we do not differentiate between relations with the same subject and those with a different subject. However we plan to conduct further experiments in the future that make use of this distinction.

In addition, we have collected *action-effect* relation instances for a baseline measure. The baseline consists of instances that cooccur with eleven patterns that indicate *action-effect* relation. The difference between the extended Espresso and baseline is caused by the low number and constant scores of patterns.

5.2 Results

We ran the extended Espresso algorithm starting with 971 positive and 1069 negative seed relation instances for *action-effect* relation and 860 positive and 74 negative seed relations for *action-means* relation. As a result, we obtained 34,993 cooccurrence patterns with 173,806 relation instances for the *action-effect* relation and 23,281 coocurrence relations with 237,476 relation instances for the *action-means* relation after 20 iterations of pattern ranking/selection and instance ranking/selection. The threshold parameters for selecting patterns and instances were decided in a preliminary trial. Some of the acquired patterns and instances for the *action-effect* relation are shown in Table 1.

5.2.1 Precision

To estimate precision, 100 relation instances were randomly sampled from each of four sections of the ranks of the acquired instances for each of the two relations (1–500, 501–1500, 1501–3500 and 3500–7500), and the correctness of each sampled instance was judged by two graduate students (i.e. 800 relation instances in total were judged).

Note that in these experiments we asked the assessors to be lenient to both (a) the degree of the likeliness that the effect/means takes place and (b) which arguments are shared between the two events. For example, while *nomu* (drink) does not necessarily result in *futsukayoi-ni naru* (have a hangover), the assessors judged this pair correct because one can at least say that the latter *sometimes* happens *as a result of* the former. For criterion (b), as shown in Table 1, the relation instances judged correct include both the *X-ga* VP_1*::X-ga* VP_2 type (i.e. two subjects are shared) and the *X-o* VP_1*::X-ga* VP_2 type (the object of the former and the subject of the latter are shared). The issue of how to control patterns of argument sharing is left for future work.

As a result of this leniency in judgement, the inter-assessor agreement turned out to be moderate. The κ statistics was 0.53 for the *action-effect* relations, 0.49 for the *action-effect* relations (=baseline) and 0.55 for *action-means*. The precision for the agreed samples are shown in Figure 1, Figure 2, Figure 3.

"2 judges" means that an instance is acceptable to both judges. "1 judges" means that it is an acceptable instance to at least one of the two judges. "strict" indicates correct instance relations while "lenient" indicates correct instance relations, when a judge append the right cases.

The figures show that both types of relations were acquired with a reasonable precision not only for the highly ranked instances but also for those close to the cut-off threshold. It may seem strange that the precision of the lower-ranked *action-means* instances is even better than the higher-ranked ones, which may mean that the scoring function given in Section 4 did not work properly. While further investigation is clearly needed, it should also be noted that higher-ranked instances tended to be more specific than lower-ranked ones and are less likely to be judged leniently.

5.2.2 Effects of seed number

We reran the extended Espresso algorithm after 20 iterations of pattern ranking/selection and instance ranking/selection to confirm effects of the seed number, starting with 500 positive and 500 negative seed relation instances about half seed instances for *action-effect* relation. The precision for the agreed samples is shown in Figure 4[2].

This precision is lower than that of *action-effect* relations with all seed instances. Additionally, the number of seed instance affects the precision of both higher-ranked and lower-ranked instances.

5.2.3 Effects of using verbal nouns

We confirmed effects of using verbal nouns. Of the 500 highest scoring patterns for acquiring *action-effect* instances that we selected, 128 include verbal nouns. For *action-means*, 495 patterns include verbal nouns. Hence, the presence of verbal nouns greatly effects some acquired instances. Additionally, of the 500 high frequent patterns selected from the 2000 highest scoring patterns for acquiring *action-effect* instances, 177 include verbal nouns, and for *action-means*, 407 patterns. This result provides further evidence that the inclusion of verbal nouns has a very positive effect in this task.

5.2.4 Argument selection

According to our further investigation on argument selection, 49 instances (12%) of the correct *action-effect* relation instances that are judged correct have a specific argument in at least one event, and all of them would be judged incorrect (i.e. over-generalized) if they did not have those arguments (Recall the example of *kogeme-ga tsuku* (get brown) in Section 4.2). This figure indicates that our method for argument selection works to a reasonable degree.

However, clearly there is still much room for improvement. According to our investigation, up to 26% of the instances that are judged incorrect could be saved if appropriate arguments were selected. For example, *X-ga taberu* (X eats) and *X-ga shinu* (X dies) would constitute an *action-effect* relation if the former event took such an argument as *dokukinoko-o* (toadstool-ACC). The overall precision could be boosted if an effective method for argument selection method were devised.

6 Conclusion and future work

In this paper, we have addressed the issue of how to learn lexico-syntactic patterns useful for acquiring event relation knowledge from a large corpus, and proposed several extensions to a state-of-the-art

[2] It was only judged by one assessor.

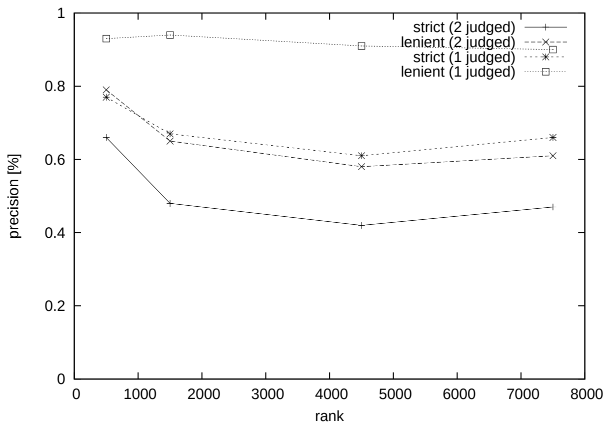

Figure 1: *action-effect*

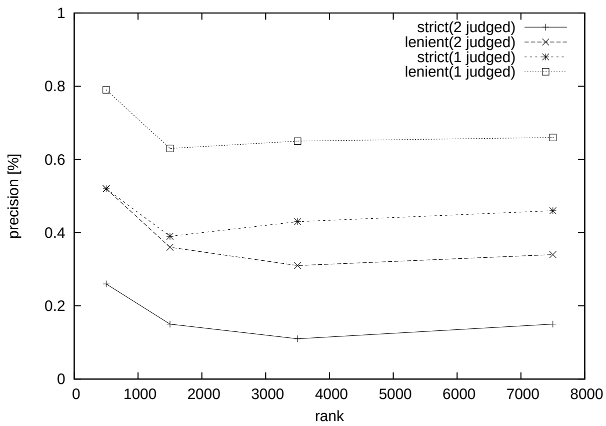

Figure 2: *action-means*

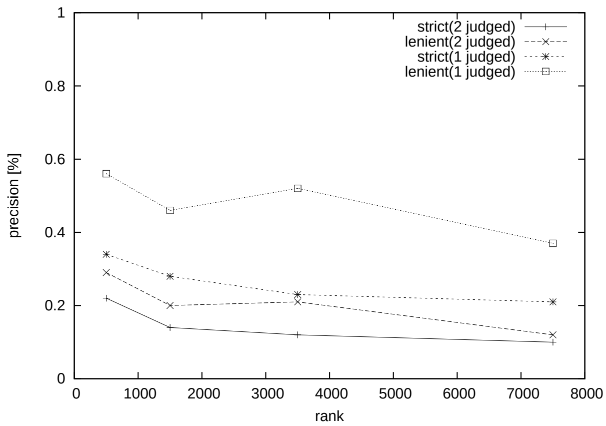

Figure 3: *action-effect* (baseline)

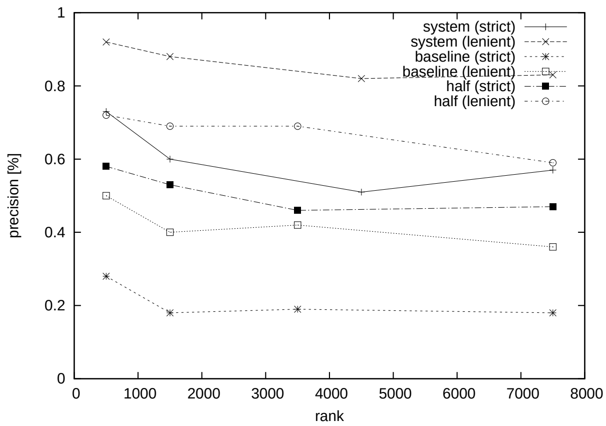

Figure 4: *action-effect* (half seed)

method originally designed for entity relation extraction, reporting on the present results of our empirical evaluation. The results have shown that (a) there are indeed specific cooccurrence patterns useful for event relation acquisition, (b) the use of cooccurrence samples involving verbal nouns has positive impacts on both recall and precision, and (c) over five thousand relation instances are acquired from the 500M-sentence Web corpus with a precision of about 66% for *action-effect* relations.

Clearly, there is still much room for exploration and improvement. First of all, more comprehensive evaluations need to be done. For example, the acquired relations should be evaluated with a strict criterion as well as the present lenient one. A deep error analysis is also needed. Second, the experiments have revealed that one major problem to challenge is how to optimize argument selection. We are seeking a way to incorporate a probabilistic model of predicate-argument cooccurrences into the ranking function for relation instances. Related to this issue, it is also crucial to devise a method for controlling argument sharing patterns. One possible approach is to employ state-of-the-art techniques for coreference and zero-anaphora resolution (Iida et al., 2006, etc.) in preprocessing cooccurrence samples.

References

Timothy Chklovski and Patrick Pantel. 2005. Global path-based refinement of noisy graphs applied to verb semantics. In *In Proceedings of Joint Conference on Natural Language Processing (IJCNLP-05)*.

Ryu Iida, Kentaro Inui, and Yuji Matsumoto. 2006. Exploiting syntactic patterns as clues in zero-anaphora resolution. In *ACL '06: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL*, pages 625–632. Association for Computational Linguistics.

Takashi Inui, Kentaro Inui, and Yuji Matsumoto. 2003. What kinds and amounts of causal knowledge can be acquired from text by using connective markers as clues? In *Proceedings of the 6th International Conference on Discovery Science,*, pages 180–193. An extended version: Takashi Inui, Kentaro Inui, and Yuji Matsumoto (2005). Acquiring causal knowledge from text using the connective marker tame. *ACM Transactions on Asian Language Information Processing (TALIP), 4(4):435–474.*

Daisuke Kawahara and Sadao Kurohashi. 2006. A fully-lexicalized probabilistic model for japanese syntactic and case structure analysis. In *In Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL2006)*, pages 176–183.

Taku Kudo and Yuji Matsumoto. 2002. Japanese dependency analysis using cascaded chunking. In *CoNLL 2002: Proceedings of the 6th Conference on Natural Language Learning 2002 (COLING 2002 Post-Conference Workshops)*, pages 63–69.

Dekang Lin and Patrick Pantel. 2001. DIRT - discovery of inference rules from text. In *In Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2001.*

Suguru Matsuyoshi, Satoshi Sato, and Takehito Utsuro. 2006. Compilation of a dictionary of japanese functional expressions with hierarchical organization. In *ICCPOL*, pages 395–402.

Patric Pantel and Marco Pennacchiotti. 2006. Espresso: Leveraging generic patterns for automatically harvesting semantic relations. In *the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the ACL*, pages 113–120.

Viktor Pekar. 2006. Acquisition of verb entailment from text. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 49–56.

Deepak Ravichandran and Eduard Hovy. 2002. Learning surface text patterns for a question answering system. In *Proceedings of the 21st International Conference on Computational Linguistics and 40th Annual Meeting of the Association for Computational Linguistics*, pages 41–47. Association for Computational Linguistics.

Kentaro Torisawa. 2006. Acquiring inference rules with temporal constraints by using japanese coordinated sentences and noun-verb co-occurrences. In *In Proceedings of Human Language Technology Conference/North American chapter of the Association for Computational Linguistics annual meeting (HLT-NAACL06)*, pages 57–64.

Fabio Massimo Zanzotto, Marco Pennacchiotti, and Maria Teresa Pazienza. 2006. Discovering asymmetric entailment relations between verbs using selectional preferences. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 849–856, Sydney, Australia, July. Association for Computational Linguistics.