

An Automatic Method of Creating Valency Entries using Plain Bilingual Dictionaries

Sanae Fujita Francis Bond
NTT Machine Translation Research Group
NTT Communication Science Laboratories
2-4 Hikari-dai Seika-cho, Kyoto, Japan 619-0237
{sanae, bond}@cslab.kecl.ntt.co.jp

Abstract

In this paper we introduce a fully automatic method to extend an existing rich bilingual valency dictionary by using information from multiple plain bilingual dictionaries. We evaluate our method using a translation regression test, and get an improvement of 7%.

1 Introduction

As the sophistication of NLP parsing increases, lexical resources with detailed syntactic and semantic information become increasingly important. However, it is expensive to build such resources. On the other hand, it is becoming increasingly easy to find large numbers of bilingual dictionaries (e.g. Sèrasset & Mangeot (2001)). In this paper we introduce a fully automatic method to extend an existing bilingual valency dictionary (hereafter, “valency dictionary”) by using information from multiple plain bilingual dictionaries (hereafter, “plain dictionary”). We extend a valency dictionary with detailed information about the argument properties of verbs and adjectives, including not just subcategorization but also selectional restrictions and translation equivalency. A plain dictionary, on the other hand, contains only translation data. The core idea is to extend the valency dictionary by creating new entries from existing entries with similar meanings. We find entries with similar meanings by comparing translations in different languages — if two Japanese words share translations in two or more different languages, then we assume that the meaning is similar.

For example, consider the case where we want to make an entry for the verb 纏う *matou* “wear”. We use a plain dictionary to find the translation *wear*. Matching through the translation gives 15 candidate Japanese verbs in the valency dictionary on which we can base the new entry. These include 着る *kiru* “wear”, 弱る *yowaru* “wear [out]”, 笑いを浮かべる *warai-o ukaberu* “wear a smile” and so on. This variety of candidates comes from the polysemy of the English verb: 着る *kiru* “wear” corresponds to WordNet sense 1 “be dressed in”, 弱る *yowaru* “wear out” to sense 8 “exhaust or tire though overuse or great strain or stress” and 笑いを浮かべる *warai-o ukaberu* “wear a smile” to sense 3 “wear an expression of one’s attitude or personality”. However, 纏う *matou* “wear” corresponds only to sense 1. We resolve this ambiguity by using a plain dictionary involving another language, such as Chinese. Because Chinese verbs have different patterns of polysemy to English, only the appropriate Japanese candidate is linked by both English and Chinese, as shown in Figure 1. The approach is similar to that of Bond et al. (2001), who create a Japanese-Malay

dictionary through linking Japanese and English through multiple pivots — the major difference is that we are not creating a new dictionary but instead increasing the size of an existing dictionary.

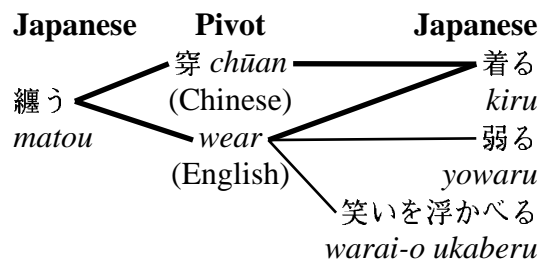


Figure 1: Determining Similarity by Linking Through Multiple Pivots

In the following, we describe the method used to create new entries in more detail (§ 2), and introduce the resources we use (§ 3). We then present different strategies for creating new valency entries based on the proposed method (§ 4), evaluate the results (§ 5) and finally discuss their implications (§ 6).

2 Method of Creating New Valency Entries

Our method takes an existing valency dictionary (the seed dictionary) and extends it automatically using plain dictionaries. Because creating valency dictionaries is more expensive than creating plain dictionaries, the coverage of plain dictionaries is generally higher. New valency entries are created by exploiting the fact that verbs with similar meanings typically have similar valency structures. That is, if there is an unknown verb in the source language (S_U) whose meaning is similar to an existing verb in the seed dictionary (the known verb S_K), we can copy the valency information of S_K for S_U .

The definition of “similar meaning” used to generate new entries is that they have the same translation. Considering translations in only one language, this massively overgenerates: one sense of a verb may overlap, but not all will. To constrain the polysemy, we propose the use of translations in multiple languages. That is, we consider two words to be similar only if they have identical translations in two or more languages.

We show the overview of our method in Figure 2, where a seed dictionary of valency pairs ($S - T$) is being extended to the new word S_U , that has no entry in the seed dictionary. In addition to the seed dictionary, we require a plain $S - T$ dictionary, larger than the seed dictionary, and one or more additional plain $S - X$ dictionaries.

Due to differences in dictionary formatting, the criteria used to match translations must be quite loose. We consider two entries the same if they have the same head word. This allows for minor inconsistencies in the target language dictionaries. Dictionaries for natural language processing often include commonly appearing adjuncts and complements that do not normally appear in plain dictionaries. Consider the case where Japanese is the source and English is the target language: 行 \langle *iku* “go” is translated as *to go* in EDICT (Breen 1995), *go* in ALT-J/E’s plain dictionary (Ikehara et al. 1991) and NP_1 *go from* NP_2 *to* NP_3 in ALT-J/E’s valency dictionary (among other transla-

Step 1: For each entry (S_U-T_U) in the plain $S-T$ dictionary with no entry in the valency dictionary

- For each valency entry (S_K) with the same target translation (T_{UK})
 - Create a candidate pair S_U-S_K

Step 2: For each candidate pair S_U-S_K (linked by T_{UK})

- Look up both source entries in a non-target plain dictionary ($S-X$)
 - Replace S_K by S_U then create a new entry (S_U-T_{UK}) for pairs which have at least one translation (X_{UK}) in common

Figure 2: Creating New Entries

tions)¹. The only word that matches is the head. Fujita & Bond (2002) examined various variations in matching other elements apart from the head, but found no improvement. So we determined the head using information in the valency dictionary in Step 1. Because of this there was additional overgeneration for complex verbs such as *give up* and *give back*.

To reduce the overgeneration, Fujita & Bond (2002) and Bond & Fujita (2003) used human judgement. But in this paper, we automatically filter out incorrect entries using matching in other languages (Step 2). Because we match the entire translation in language X , there is no overgeneration due to complex verbs. If there are multiple plain dictionaries, then the criterion in Step 2 can be varied further — for example to use all dictionaries and select these words which have at least one matching translation of X (we call this **UNION**) or to use all dictionaries and select only those words which have matching translations in all languages (we call this **INTER**).

3 Resources

In this section we describe the resources used in our experiment. The source language is Japanese, and the target language English. We use additional plain Japanese- X dictionaries for three languages: Chinese, French and German.

3.1 The Seed Valency Dictionary

We use the valency dictionary from the Japanese-to-English machine translation system **ALT-J/E** (Ikehara et al. 1991) as a seed dictionary. It is a large hand-built dictionary, primarily built for machine translation, but has been used for a variety of tasks, such as zero-pronoun detection and resolution (Yamura-Takei et al. 2002) and paraphrasing (Takahashi et al. 2001). The dictionary includes valency (subcategorization) information for Japanese verbs and their English translations, as well as selectional restrictions on the arguments. Constructing entries is expensive, taking an expert lexicographer 30 minutes. A simplified entry for 行く *iku* “go” is shown in Figure 3.

¹**ALT-J/E** is a Japanese-to-English machine translation system which has both plain dictionary and valency dictionary

Japanese Side:		English Side:
┌ N1	<agents, animal, vehicles> "が" <i>ga</i>	┌ SUBJ N1
└ N3	<creation (others),	└ VERB "go"
	places, place > "に/へ/まで" <i>ni/e/made</i>	└ PP "from" N3 OBJ-form
└ N4	<places, place> "から/より" <i>kara/iori</i>	└ PP "to" N4 OBJ-form
└ VERB	"行く" <i>iku</i>	

Figure 3: 行く *iku* "N1 go from N4 to N3"

For the non-idiomatic part of **ALT-J/E**'s valency dictionary, on the Japanese side, there are 5,134 types of verb giving 11,310 entries and 1,339 types of adjective giving 2,201 entries.² On the English side, there are 2,421 types of verb for 13,549 entries and 1,513 types of adjective for 4,664 entries (some entries have both verb and adjective). In general, translation tends to simplify text, because the target language will not be able to represent the exact same shades of meaning as the source text: therefore, the English variation is less than the Japanese variation in this dictionary.

To test the coverage of the Japanese verbs in a newspaper, we investigated the coverage of verbs in one year's worth of the Japanese newspaper articles (Nihon Keizai Shimbun 1995). Only 43% of verb types are covered, but the coverage of tokens is around 93%. Note, that this is a very simple check that considers only whether an entry exists — not whether it has an appropriate subcategorization or translation. To test the coverage of the English verbs in a newspaper, we investigated the coverage of verbs in one year's worth of the Wall Street Journal (1996), POS tagged with the Brill tagger (Brill 1995). Only 23% of verb types are covered, but the coverage of tokens is around 93%. However, as for Japanese, we are not checking whether the entries have appropriate subcategorization or translations, or considering verb-particle constructions, so this is only an upper bound of the coverage.

3.2 The Plain Bilingual Dictionaries

To create candidate entries, we use two plain Japanese-English dictionaries: **ALT-J/E**'s word transfer dictionary (Ikehara et al. 1991) and EDICT (Breen 1995). These have wider coverage than the seed dictionary, although with less detailed information. To constrain polysemy we used a Japanese-to-Chinese machine dictionary available in machine readable form: *J - C* (Shogakukan & Peking Shomoinshokan 1987); and two dictionaries available on-line: Wadoku Jiten — a Japanese-to-German dictionary *J - G* (Apel 2002); and Dico FJ — a Japanese-to-French dictionary *J - F* (Desperrier 2002). In Table 1, we show the number of entries for each of the plain dictionaries used in this paper. Three of these dictionaries (EDICT, Wadoku Jiten and Dico FJ) are available on-line, and are growing over time; the numbers given here are for the versions we used. Most bilingual entries lacked POS tags, so we matched on the surface form of all entries, even though most are not verbs or adjectives.

²**ALT-J/E**'s valency dictionary has both a common and an idiomatic structure transfer dictionary.

Table 1: Size of J-X Dictionaries

J-X	Japanese	X	Pairs
J-C	72,400	102,300	180,800
J-G	252,400	224,000	526,000
J-F	16,600	10,500	37,900
J-E (EDICT)	94,200	80,400	154,600
J-E (ALT-J/E)	323,700	276,100	415,000

Table 2: No. of Created Valency Entries

Strategy	Created Entries		Verbs	
	No.	%	No.	%
CN	2,077	22.6	668	35.8
DE	7,826	85.3	1,694	90.7
FR	629	6.9	153	8.2
INTER	141	1.5	51	2.7
UNION	9,178	100.0	1,868	100.0

4 Valency Entry Creation

We use the plain dictionaries in several ways: That is, we only use the pairs of S_U and S_K which have the same: (1) Chinese translation C (we call this strategy **CN**), (2) German translation G (we call this **DE**), (3) French translation F (we call this **FR**), (4) have at least one matching translation in C , G and F (**UNION**), or (5) have matching translations in all of C , G and F (**INTER**).

Table 2 shows the number of valency entries created by each strategy. As **UNION** creates the most entries, the number created by each method is also shown as a percentage of this total. The number of candidates created using only English data (without a pivot language) is 132,111 for 4,129 verbs: most of these are spurious entries.

5 Translation-based Evaluation

5.1 Method

We evaluated our method using a translation regression test, using the Japanese-to-English machine translation system **ALT-J/E**. We translated the test sentences both **with** the valency dictionary which has the new entries, and **w/out** new entries. When there is no entry for a verb in the valency dictionary, the system uses either the default translation in the plain dictionary or if there is no entry in the plain dictionary, the Japanese verb as is.

Translations that were identical were marked **no change**. Translations that changed were shown to evaluators with the **with** and **w/out** translations placed in random order. For example, in (1)³ A is **w/out** and B is **with**.

- (1) 動くものがいると心がなごむものです。
 ugoku mono ga iru to kokoro ga nagomu mono desu.
 move thing NOM exist if heart NOM calm down which is
 A: If there is a thing which moves, a heart is softened.
 B: If there is a thing which moves, we calm down.

³We use the following abbreviations: NOM: nominative postposition; DAT: dative postposition; LOC: locative postposition; ACC: accusative postposition;

Translations that changed were placed into three categories: (i) A is **better** than B, (ii) A is **equivalent** in quality to B and (iii) A is **worse** than B. In (1), the change is judged to be (iii): A is **worse** than B. As it happens, A is **w/out** and B is **with**: that is **with** is **better** than **w/out**.

5.2 Results Over Common Test Set

To compare the 5 strategies, we consider the test set made up of sentences with results for all five strategies. Of course, the number of verbs created by **INTER** is the smallest. We tried to get two sentences for each verb S_U created by **INTER**, but could only find one sentence for some verbs: this gave a total of 101 test sentences. The test sentences were extracted from Japanese newspapers.

Translations that changed were evaluated by 3 Japanese native speakers whose TOEIC scores are higher than 910, on average 927.⁴ Each of the three evaluators evaluated all sentences. To see the concordance rate of the evaluations by the 3 evaluators, we used Kendall's coefficient of concordance W (Muto 1995). For the calculation, we set **better** to 3, **eq** to 2 and **worse** to 1. Kendall's coefficient of concordance W is 0.57 and the P value is $P = 4.3 \times 10^{-19}$, this indicates that the data is significant and agreement between the three evaluators is high.

We summed the results of the three annotators into a single score (*score*). If the evaluation is **better**, we add 1 to *score*, and if **worse**, we subtract 1 from *score*. The value of *score* thus ranges from -3 to 3. Figure 4 shows the results. We discuss the results in Section 6.1.




Figure 4: Comparative Evaluation of Impact on Translation Quality

5.3 Overall Results

In Section 5.2, we compared several strategies, using the same test sentences for all the strategies. In this section, to evaluate all the created entries, we selected the test sentences for all S_U created by each strategy. In this case each sentence was evaluated by a single evaluator. Table 3 shows the results. The overall improvement (**improved** – **degraded**) ranges from 4.7% to 8.7%. The

⁴TOEIC (Test of English for International Communication) is a standardized test of English ability, with scores ranging from 10 to 990. <http://www.toeic.com/>

Table 3: Overall Evaluation of Impact on Translation Quality for each Strategy

Impact on Translation Quality	Strategy							
	CN		DE		FR		UNION	
	No.	%	No.	%	No.	%	No.	%
improved	305	23.5	776	24.7	39	16.6	873	24.2
equivalent	392	30.2	991	31.6	98	41.7	1,153	31.9
no change	410	31.6	809	25.8	70	29.8	950	26.3
degraded	192	14.8	561	17.9	28	11.9	634	17.6
difference	113	+8.7	215	+6.8	11	+4.7	239	+6.6
Total	1,299	100.0	3,137	100.0	235	100.0	3,610	100.0

biggest improvement was for **CN**, which comes from a totally different language family to English. In all cases, the number of **improved** sentences is greater than those **degraded**, but **UNION** creates the most entries, and has an overall improvement of 6.6%.

6 Discussion and Further Work

6.1 Comparison of the Strategies

As we can see in Figure 4, the 5 strategies have comparative performance. **UNION**, which is the least restrictive condition, has more **improved** sentences than **INTER**, which is the strictest condition. And **UNION** has fewer **degraded** sentences than any one language used individually.

We hypothesize that the reason for **UNION**'s strength is that it creates entries for different senses. Because the entries include selectional restrictions, the machine translation system is able to select the most suitable entry. (2) is an example where the translation of **UNION** is improved, but **FR** is degraded.

- (2) その悔しさを晴らすべく、2国はフランスに乗り込んでくる。
 sono kuyashisa o harasu beku, 2-koku wa France ni norikonde kuru.
 the chagrin ACC dispel to, 2 nations NOM France DAT enter come.
 w/out: It should dispel that chagrin and two countries get into in France.
 with(**FR**): It should clear that chagrin and two countries enter France.
 with(**UNION**): It should dispel that chagrin and two nations enter France.

In (2), there are two changes between **w/out** and **with**, which are underlined. Indeed, according to Gakken (Kindaichi & Ikeda 1988), 晴らす *harasu* has two senses: *clear* and *dispel*. **FR** made only the one entry: *clear*; but **UNION** made entries for both *clear* and *dispel*, so the system could select the more suitable entry. Further, both **FR** and **UNION** translated 乗り込む *norikomu* as *enter*, but **UNION** had more informative selectional restrictions, so the translation of 2国 *2-koku* “two nations” became the more suitable *two nations*.

(3) is an example where **UNION** had a bad effect.

- (3) 京都 駅 で 降りる。
 Kyoto eki de oriru.
 Kyoto station LOC get off.
 w/out: * get down in Kyoto Station.
 with(**UNION**): It withdraws in Kyoto Station.
 with(**DE**): It falls in Kyoto Station.
 with(**INTER**): It lands in Kyoto Station.

降りる *oriru* has many senses. According to Gakken (Kindaichi & Ikeda 1988), it has 10 senses including *get off/land, withdraw, fall*. In (3), 降りる *oriru* is used with the meaning *get off*, so **UNION** and **DE** are somewhat degraded. However the created entries are not bad, just inappropriate, so if more suitable entries were created, this degradation could be resolved.

The results show that the **UNION** strategy makes few wrong entries and the quality is almost as high as **INTER**. In addition, **UNION** creates entries for far more verbs (1,868) than **INTER** (51) (see Table 2). Overall, the quality of sentences translated using these verbs increases by 6.7%. Adding all the automatically constructed entries increased the token coverage over the Japanese newspaper text from 93% to 96%, while the verb types coverage from 48% to 76%.

In summary, **UNION** is the best strategy when there are multiple plain dictionaries available.

6.2 Analysis of Translation Variation

In this section, we analyse the reasons for the **improved** and **degraded** translations.

There were three main reasons why translations **improved**: (i) a new entry with a better translation was selected (see (1) or **UNION** of (2)); (ii) the sentence was translated using the correct subcategorization, which allowed a zero pronoun to be supplemented or some other improvement; and (iii) the system was able to translate a previously unknown word.

Next, we discuss the sentences which **degraded**. The main reason is that the translation of the verb became less appropriate. Often, the entry created was good, but not appropriate for the sentence being translated. Occasionally we also created a more appropriate entry, but the selectional restrictions did not enable the translation system to choose it. A secondary reason is that the new entry caused resolving zero pronoun resolution to fail, especially for long sentences.

For example, 飲み込む *nomikomu* has the core sense of *swallow* which is extended metaphorically to mean *understand*. However, the seed dictionary did not have *swallow* in the core sense “pass through the esophagus as part of eating or drinking” and so *nomikomu=understand* was the only entry produced. Therefore, any sentences where the original sense was appropriate were mistranslated. However, the valency dictionary didn’t originally contain an entry for 飲み込む *nomikomu*, causing **ALT-J/E** to use the default translation of *swallow*.

(3) is an example where the appropriate entry was created, but the system did not select it. For 降りる *oriru* “get off/down”, the system created many entries: *get out, land, withdraw* and *fall*. However, the translation system has trouble selecting the appropriate translation. To make the current system select the most suitable entry, the selectional restrictions must be tuned. But ideally,

if the Japanese side has no real sense distinctions, the selection of the translation should be done on the English side.

For example, in (4) and (5), for the verb 降りる *oriru*, the Japanese side doesn't need different valency information. However, the English side needs different valency information due to the difference between *get out [of the taxi]* and *get off [the ferry]*. This is a problem with the underlying translation system.

- (4) 神戸 で タクシー を 降りる。
Kobe de taxi o oriru.
Kobe at taxi ACC get out of
I get out of the taxi at Kobe.
- (5) 神戸 で フェリー を 降りる。
Kobe de ferry o oriru.
Kobe at ferry ACC get off
I get off the ferry at Kobe.

Examining the entries, in general, the entries which cause the degradation are not bad entries. Even the **UNION** strategy uses the intersection of J-E and at least one J-X, so few spurious entries are created. The main cause of the degradation is the fact that we have not created enough entries to cover all the senses. So, we need to add more entries which can cover all of the senses. That is, to improve the translation, we need a seed dictionary with a greater coverage of variation in the target language. In addition, if the plain dictionaries were richer, we could both match with higher precision and exploit more information in making the subcategorization frames.

6.3 Comparison with Other Approaches

Our work differs from corpus-based work such as Kawahara & Kurohashi (2001) in that we are using existing lexical resources rather than a corpus. This gives an advantage with low frequency words, so long as we can find them in a plain dictionary. It also allows us to create bilingual entries. The method gives less improvement than semi-automatic approaches such as Bond & Fujita (2003), who claim an improvement of around 16%. Using human judgments to restrict the candidates allows more entries to be created, but at a cost of 6 minutes per entry. The advantage of the approach proposed in this paper is that it is fully automatic, so it can be redone whenever a new bilingual dictionary is available.

6.4 Further Work

In this paper we used Japanese as the source and English as the target. However, the method can be used in either direction. We next plan to go the other way, using English as the source and Japanese as the target. This will allow us to flesh out the number of different English translations, as well as taking advantage of the huge numbers of plain English-X dictionaries. In addition we want to try this approach with existing entries, to see if we can add new translations. Once we have more variety on the English side, we will investigate extending the dictionary once more with English as the target.

7 Conclusion

We proposed a fully automatic method to expand an existing valency dictionary which has rich information, using plain bilingual dictionaries from multiple language pairs. We showed that translations could be improved in 7% of sentences by creating new entries based on existing entries with the same English translation and one or more identical translations in a third language. In this way, fully exploiting existing resources helps us to break through the knowledge acquisition bottleneck.

References

- Apel, Ulrich: 2002, 'WaDokuJT - A Japanese-German Dictionary Database', in *Proceedings of Papillon 2002 Workshop (CDROM)*, (Lexicon at: <http://www.babbletower.net/>).
- Bond, Francis & Sanae Fujita: 2003, 'Evaluation of a method of creating new valency entries', in *MT Summit IX*, New Orleans, pp. 16–23.
- Bond, Francis, Ruhaida Binti Sulong, Takefumi Yamazaki & Kentaro Ogura: 2001, 'Design and construction of a machine-tractable Japanese-Malay dictionary', in *MT Summit VIII*, Santiago de Compostela, Spain, pp. 53–58.
- Breen, Jim: 1995, 'Building an electronic Japanese-English dictionary', Japanese Studies Association of Australia Conference (http://www.csse.monash.edu.au/~jwb/jsaa_paper/hpaper.html).
- Brill, Eric: 1995, 'Transformation-based error-driven learning and natural language processing: A case study in part-of-speech tagging', *Computational Linguistics*, **21**: 543–65.
- Desperrier, Jean-Marc: 2002, 'Analyze of the results of a collaborative project for the creation of a Japanese-French dictionary', in *Proceedings of Papillon 2002 Workshop (CDROM)*, (Lexicon at: <http://dico.fj.free.fr/dico.php>).
- Fujita, Sanae & Francis Bond: 2002, 'Extending the coverage of a valency dictionary', in *COLING-2002 workshop on Machine Translation in Asia*, Taipei, pp. 67–73.
- Ikehara, Satoru, Satoshi Shirai, Akio Yokoo & Hiromi Nakaiwa: 1991, 'Toward an MT system without pre-editing – effects of new methods in **ALT-J/E**', in *Third Machine Translation Summit: MT Summit III*, Washington DC, pp. 101–106.
- Kawahara, Daisuke & Sadao Kurohashi: 2001, 'Japanese case frame construction by coupling the verb and its closest case component', in *Proceedings of First International Conference on Human Language Technology Research (HLT 2001)*, San Diego, pp. 204–210.
- Kindaichi, Haruhiko & Yasaburo Ikeda: 1988, *Gakken Japanese Dictionary 2nd edition*, Gakken Co., Ltd.
- Muto, Shinsuke: 1995, *Toukei Kaiseki no Handbook [Handbook of Statistic Analysis]*, Asakura Shoten, (in Japanese).
- Sèrasset, Gilles & Mathieu Mangeot: 2001, 'Papillon lexical database project: Monolingual dictionaries & interlingual links', in *Proceedings of the Sixth Natural Language Processing Pacific Rim Symposium: NLPRS-2001*, Tokyo, pp. 119–125.
- Shogakukan & Peking Shomoinshokan, eds.: 1987, *Ri-Zhong Cidian [Japanese-Chinese Dictionary]*, Shogakukan.
- Takahashi, Tetsuro, Tomoya Iwakura, Ryu Iida, Atsushi Fujita & Kentaro Inui: 2001, 'KURA: a transfer-based lexico-structural paraphrasing engine', in *Proceedings of the 6th Natural Language Processing Pacific Rim Symposium (NLPRS) Workshop on Automatic Paraphrasing: Theories and Applications*, pp. 37–46, <http://cl.aist-nara.ac.jp/kura/doc/>.
- Yamura-Takei, Mitsuko, Miho Fujiwara, Makoto Yoshie & Teruaki Aizawa: 2002, 'Automatic linguistic analysis for language teachers: The case of zeros', in *19th International Conference on Computational Linguistics: COLING-2002*, Taipei, pp. 1114–1120.