

Enhancing a Dictionary for Transfer Rule Acquisition

Francis Bond*, Eric Nichols,** and Jim Breen***

* NICT Computational Linguistics Group

** Nara Institute of Science and Technology

*** Monash University

<bond@ieee.org,eric-n@is.naist.jp,
jim.breen@infotech.monash.edu.au>

Abstract

The JMdict/EDICT Japanese-English Dictionary is a freely-available dictionary distributed in XML (JMdict) and text (EDICT) formats. It is widely used as a source of lexical material in dictionary systems and text-processing projects. We propose two refinements to make the dictionary more computationally tractable: marking entries where the English is not a translation equivalent and expanding contracted entries. We then propose and apply semi-automatic methods to refine existing entries. The resulting dictionary is shown to be more suitable for the construction of machine translation rules.

1 Introduction

Resources built for one task can often be useful in others. WordNet, for example, started off as a test-bed for a particular model of lexical organization (Fellbaum, 1998, p4) and is now widely used in natural language processing (NLP) applications. In this paper we look at the Japanese/English lexicon JMdict/EDICT (Breen, 2004), which started out as a voluntary project to produce a freely available Japanese/English Dictionary in machine-readable form. In addition to being useful for people as a bilingual dictionary, it is also widely used in NLP applications. For example, it has been the base to make compound noun lexicons (Tanaka and Matsuo, 1999; Ohmori and Higashida, 1999), new bilingual lexicons (Paik et al., 2001; Apel, 2002; Sjöbergh, 2005; Zhang et al., 2005; Fujita and Bond, 2006; Bond and Ogura, 2007) and machine translation transfer rules (Bond et al., 2005; Nichols et al., 2007).

We look at two ways of making the dictionary even more tractable for NLP tasks: (1) marking entries where the English is not a translation equivalent and (2) expanding contracted entries. Then, we discuss some planned future enhancements. Finally we describe the process of making transfer rules in more detail.

2 JMdict/EDICT

The JMdict/EDICT project now has approximately 110,000 Japanese/English entries recorded, with the number increasing at about 1,000 per month. A WWW-based system for submitting amendment and new entry suggestions is yielding about 100 submissions each day, which close to the limit that can be handled by the sole editor (Breen).

The project dictionary is distributed in three formats:

1. the full JMdict (XML) format, both in Japanese-English and Japanese-English/German/French/etc. versions.
2. the original EDICT format, which only allows for one kanji word and one reading per entry. Thus JMdict entries which have alternative kanji or okurigana forms, or which have alternative readings will result in multiple entries in the EDICT file.
3. the EDICT2 format (shown below) which allows for multiple kanji words and readings in an entry, and is in effect a human-readable equivalent of the JMdict entries.

The dictionary files are generated daily, and are available via ftp and rsync, thus allowing WWW servers that use the files to stay up-to-date.¹

At present all editing is taking place at Monash University, with semi-automated creation of new entries, and manual amendment of existing entries. A new WWW-based maintenance system is nearing completion which will enable distributed editing with a pool of editors. The new system has a more flexible database which will allow additional information to be included in the entries, and greater access to the data by project members.

3 Enhancements to the Dictionary Structure

A typical entry (for *jiten* “dictionary”) in the EDICT2 format is shown below:

(1) 辞典(P);辭典(oK) [じてん] /(n) dictionary/(P)/

The various marks indicate that it is a common word (P), there is an orthographic variant with old Kanji (oK), and that the Japanese part of speech is a noun (n). In this entry, as in most entries, the English gloss is a translation equivalent and the entry is effectively reversible: 〈辞典 ↔ dictionary〉. This allows the use of JMdict as an English→Japanese lexicon, even though it basically Japanese→English.

However there are some entries where the reversibility does not hold. For example consider the simplified entry for *ten* “piece” (2):

¹<http://ftp.monash.edu.au/pub/nihongo/00INDEX.html>

- (2) 点 [てん] (1) /(n,n-suf) spot/mark/
 (2) point/dot/
 (3) (n-suf) counter for goods or items/(P)/

In this case the third gloss is not a translation equivalent, but rather an explanation: *ten* “piece” is used as a suffix for numbers when counting goods or items. We would not expect to want to look up this directly, and could not directly use the gloss to create translation rules.

Another example where the reverse look up fails is disjunctive entries such as (2):

- (3) 田地 [でんち; でんじ] /(n) farmland/rice field or paddy/

In this case, two translations have been collapsed into the second gloss: *rice field* and *rice paddy*. This contraction of entries is important in paper dictionaries, where space is precious, but causes problems for electronic access: the translation equivalent *rice field* will not be an exact match and the translation *rice paddy* is not even a contiguous string.

The solution to the first problem is to explicitly mark the type of each gloss. The default type is `equ` (translation equivalent) whereas explanations are marked as `exp`. Simplified examples of this marked up in XML are shown in (4):

- (4) 点 [てん] ...
`<gloss g_type="equ">spot</gloss>`
`<gloss g_type="exp">counter for goods or items</gloss>`

The solution to the second is even simpler: split the entry with “or” into two separate entries:

- (5) 田地 [でんち; でんじ] /(n) farmland/ rice field/ rice paddy/

4 Expanding Disjunctive Entries

There were 2843 entries containing “or” in JMdict (1.3% of the Japanese-English entries contained a disjunctive gloss). Four word entries were the most common, with the longest entry consisting of 35 words.

An initial survey of the glosses found three major types (**G**) good translation equivalents (**D**) disjunctive glosses and (**E**) explanations.

- G** 是非か [ぜかひか] /right or wrong/
 有無 [うむ] /yes or no/

- D** 国際語 [こくさいご] /(n) an international or universal language/
 歳入 [さいにゅう] /(n) annual revenue or income/

- E** 読み破る [よみやぶる] /(v5r,vt) to read through (difficult passage or particularly long book)/

The vast majority were of short entries (5 words or less) were of type **D**, while the longer entries were mainly of type **E**.

The algorithm for rewriting was simple:

1. Remove any articles from the gloss
2. if the final word is *two*, *other* or *another*
⇒ **G**
3. elsif the Japanese entry ends in か
⇒ **G**
4. elsif the gloss appear more than 3 times (e.g *yes or no*)
⇒ **G**
5. elsif there 6 or more words
⇒ **E**
6. elsif *or* is the second word (w_2)
D split into w_1, w_4, w_5, \dots and w_3, w_4, w_5, \dots
7. elsif *or* is the second last word (w_{n-2})
D split into $\dots w_{n-4}, w_{n-3}, w_{n-2}$ and $\dots w_{n-4}^1, w_{n-3}, w_n$
8. else leave to be hand-checked

As a result of this 72 entries were judged to be good (**G**), 1,500 to be disjunctive (**D**) and the remaining 1,271 to be explanations (**E**) or requiring further checking.

In the examples given above, the disjunctive entries are rewritten as follows:

- (6) 国際語 [こくさいご] / (n) international language/ universal language/
- (7) 薄遇 [はくぐう] / (n,vs) poor reception/ inhospitable reception/

5 Evaluation

The effectiveness of splitting was tested by manual evaluation. For those judged **G**, there were only three errors, all caused by errors in the original entries.

For those split (**D**), only 2% were erroneous. The main source of errors was splitting good entries, such as (8).

- (8) 虚実 [きょじつ] / (n) truth or falsehood/

The 1,271 **E** entries included 304 that were actually disjunctive, but the vast majority were explanations such as (9), which was corrected to (10)

- (9) 押し鮨;押し寿司 [おしずし] /(n) {food} sushi rice and other ingredients pressed in box or mould (mold)/
- (10) 押し鮨;押し寿司 [おしずし] /(n) {food} oshizushi/ (expl.) sushi rice and other ingredients pressed in box or mould/

We also investigated testing the validity of splitting the disjunctive entries by looking them up in a different lexicon (EDR, 1990), but found only 35 hits, too few to be useful. Similarly, we investigated looking up entries in a bilingual aligned corpus (Utiyama and Takahashi, 2003), but found too few hits to be useful.

6 Future Enhancements of JMDict

A number of enhancements are under consideration and at various stages of implementation. They include:

1. Extension of frequency-of-use information. At present about 30% of entries have frequency ranking information based primarily on a newspaper-based ranked word-list. This is being extended and refined using WWW-based word-frequency metrics.
2. Expansion of orthographical variants, including okurigana variants and kana substitution for non-Joujou kanji. Experimentation is under way into the automatic generation of potential variants combined with testing possible variants against the WWW and other corpora to determine their validity and level of usage.
3. Greater delineation of senses. At present only about 5% of entries have senses marked in the English glosses. While Japanese is not regarded as being highly polysemous, there is considerable scope to improve the level of sense tagging. There is some potential to employ NLP techniques on the English glosses to identify candidates for sense delineation.
4. Extension of cross-referencing including indication of synonyms and antonyms. At present there is a relatively low level of cross-referencing. Some experimentation using bag-of-words techniques with the English glosses has shown that this may be a fruitful approach for identifying synonyms.
5. Marking of domains. At present there is limited domain marking in a number of entries, and it is highly desirable that this be extended. An issue is determining an appropriate set of word domains to use. A possibility being explored is the application of Wordnet synsets.

6. Adding verb translations for verbal nouns (サ変名詞). At present EDICT generally records only the noun translation: for example: 検査 [けんさ] / (n,vs) inspection/. We would like to expand the entry to: 検査 [けんさ] / (n) inspection/ (vs) inspect/. This can be done semi-automatically and hand checked, as we did for the disjunctive entries. In this case the verbal form can be deduced from the nominal one using Nomlex (Macleod et al., 1998) or WordNet 2.0. However, there are over 10,000 verbal nouns, so semi-automatic checking becomes even more important.

7 Rule Acquisition

In this section we introduce how we produce transfer rules for a machine translation system. The specific system we are targetting is JaEn, a semantic transfer system using HPSG grammars for both parsing and generation (Bond et al., 2005).

7.1 Japanese→English MT

The architecture of our Japanese→English system (hereafter referred to as “Jaen”) is semantic transfer via rewrite rules, as shown in Figure 1. The source text is parsed using an HPSG grammar for the source language, and a semantic analysis in the form of Minimal Recursion Semantics (MRS) is produced. That semantic structure is rewritten using transfer rules into an target language MRS structure, which is finally used to generate text from a target language HPSG grammar. Jaen uses the MRS rewrite translation engine from the LOGON² Norwegian→English MT system (Lønning et al., 2004) and parsers as well as English and Japanese grammars freely available from the DELPH-IN project³. In particular, we use the Japanese and English HPSG-based grammars JACY (Siegel, 2000) and the English Resource Grammar (ERG; Flickinger (2000)).

MRS (Copestake et al., 2005) is a semantic formalism representing syntactic relations in a generalized manner. MRS structures are flat, unordered collections of elementary predications (EPs) with handles (h) indicating scopal relations, events (e), and entities (x). Figure 2 gives the MRS for the sentence “Research is fun.” The sentence is a statement, and the message, `proposition_m_rel(e2)` indicates this. `tanoshii_a_rel(e2,x6)` is an event, and takes `kenkyuu_s_rel(x6)` as its subject. `noun-relation(x6)` nominalizes `kenkyuu_s_rel(x6)`, which is normally an event, turning it into an entity. MRS provides several features that make it attractive as a transfer language, such as uniform representation of pronouns, specifiers, temporal expressions,

²<http://www.emmtee.net>

³<http://www.delph-in.net>

and the like over grammars. More details can be found in Flickinger et al. (2005).

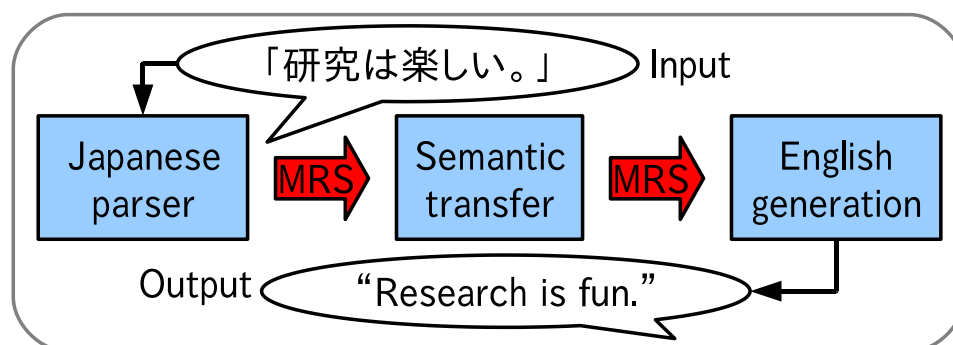


Figure 1: The Jaen machine translation architecture

7.1.1 Transfer Rules

As illustrated in Nygård et al. (2006), transfer rules take the form of MRS tuples:

```
[CONTEXT:] IN[!FILTER] -> OUT
```

where IN(PUT) is rewritten by OUT(PUT), and the optional CONTEXT specifies relations that must be present for the rule to match, and conversely, FILTER specifies relations whose presence blocks a rule from matching. Consider the following transfer rule to translate “言語” *gengo* into “language”:

```
gengo-language-mtr :=
[ IN.RELS < [ PRED"_gengo_n_1_rel",
              LBL #h1, ARG0 #x1 ] >,
  OUT.RELS < [ PRED"_language_n_1_rel",
              LBL #h1, ARG0 #x1 ] > ].
```

This rule rewrites any instance of *gengo.n.1.rel* with *language.n.1.rel*. *#h1* and *#x1* indicate that the LBL and ARG0 arguments of the MRS produced must be preserved. While this may seem like a fairly easy to understand rule, we must repeat the constraint on LBL and ARG0 every time we write a rule to translate nouns. In order to avoid such redundancy in rule writing, LOGON allows the user to specify rule types that can encapsulate common patterns in rules. The above rule can be generalized to cover nouns:

```
noun_mtr :=monotonic_mtr&
[IN.RELS < [ LBL #h1, ARG0 #x1 ] >,
  OUT.RELS < [ LBL #h1, ARG0 #x1 ] > ].
```

```

研究 が 楽しい
[ LTOP: h1
  INDEX: e2 [ e TENSE: PRES
              MOOD: INDICATIVE
              PROG: - PERF: - ]
  RELS: <
    [ PRED proposition_m_rel
      LBL: h1
      ARG0: e2
      MARG: h3 ]
    [ PRED "_kenkyuu_s_rel"
      LBL: h4
      ARG0: x5
      ARG1: u7
      ARG2: u6 ]
    [ PRED "noun-relation"
      LBL: h8
      ARG0: x5
      ARG1: h9 ]
    [ PRED proposition_m_rel
      LBL: h9
      ARG0: x5
      MARG: h10 ]
    [ PRED udef_rel
      LBL: h11
      ARG0: x5
      RSTR: h12
      BODY: h13 ]
    [ PRED "_tanoshii_a_rel"
      LBL: h14
      ARG0: e2
      ARG1: x5 ] >
  HCONS: < h3 qeq h14, h10 qeq h4,
           h12 qeq h8 > ]

```

Figure 2: MRS for *kenkyuu-ga tanoshii* "research is fun"

and our example rule can be rewritten as:

```
gengo-language-mtr := noun_mtr &
[ IN.RELS < [PRED "_gengo_n_1_rel" ] >,
  OUT.RELS <[PRED "_language_n_1_rel"]>].
```

The LOGON system contains a rich definition of rule types - many of which were immediately applicable to Jaen. Jaen inherited from LOGON rule types for open category lexical items such as common nouns, adjectives, and intransitive & transitive verbs. In addition, LOGON contains a number of rule types to specify rules for quantifiers, particles, and conjunctions, providing much of the framework needed to develop Jaen.

7.1.2 Rule Types Unique to Jaen

Here, we briefly describe a few rule types that were developed to handle linguistic phenomena unique to Japanese→English translation. In Figure 2, we see an example of the Japanese verbal noun, 研究 “research” being used as a noun. In Jaen, Japanese verbal nouns are analyzed as events, and they produce messages accordingly. When it is being used as a noun, `kenkyuu_s_rel` is wrapped with the relation `noun-relation`. We handle these constructions with a special rule that nominalizes the verbal noun by removing its event and the associated message, and replacing them with an entity when it appears as a noun:

```
vn-n_jf := monotonic_mtr &
[ CONTEXT.RELS < [ PRED "ja:undef_rel",
                  ARG0 #x0 ] >,
  IN [RELS <[PRED "ja:noun-relation",
             LBL #h6, ARG0 #x0, ARG1 #hp],
      [PRED "ja:proposition_m_rel",
          LBL #hp, ARG0 #ep, MARG #h5 ],
      [PRED #pred, LBL #h0, ARG0 #ep ]>,
    HCONS < qeq & [HARG #h5, LARG #h0 ]>],
  OUT [RELS <[PRED #pred, LBL #h6,
             ARG0 #x0 ]>,
      HCONS < > ] ].
```

In short, this rule type removes the `noun-relation` and all semantic relations resulting in the verbal noun’s analysis as an event. This change makes it possible to treat verbal nouns identically to regular nouns in the rest of our transfer rules, eliminating the need to create multi-word transfer rules that have to distinguish between nouns and verbal nouns. This simplifies rule development significantly. And, thus, a rule to translate 「研究」 as the noun “research” can now be created using the standard noun template:

```
kenkyuu_s-research_n-omtr := noun_mtr &
[IN.RELS <[PRED "_kenkyuu_s_rel"]>,
  OUT.RELS<[PRED "_research_n_1_rel"]>].
```

7.2 Acquiring Open Category Transfer Rules from Bilingual Dictionaries

Nygård et al. (2006) demonstrated that it is possible to learn transfer rules for some open category lexical items using a bilingual Norwegian→English dictionary. They succeeded in acquiring over 6,000 rules for adjectives, nouns, and various combinations thereof. Their method entailed looking up the semantic relations corresponding to words in a translation pair, and matching the results using simple pattern matching to identify compatible rule types. Our approach is an effort to generalize this approach by using rule templates to generate transfer rules from input source and target semantic structures. We use Minimal Recursion Semantics (MRS) as our semantic representation. Template mappings are used to identify translation pairs where there is a compatible rule type that can be used to create a transfer rule. A template mapping is a tuple consisting of: (i) a list of HPSG syntactic categories corresponding to the words in the source translation; (ii) a list of HPSG syntactic categories for the target translation words; and (iii) the name of the rule template that can be used to construct a transfer rule. For example, the template mapping ([noun], [adjective, noun], n-an) identifies a template that creates a rule to translate a Japanese noun into an English adjective-noun sequence. We use the enhanced version of JMDict described above. Transfer rule generation is carried out in the following manner:

1. Look up all words in source language translation in HPSG grammar
 - Retrieve syntactic categories and MRS relations
 - Enumerate every possible combination for words with multiple entries
 - Refactor results into separate lists of syntactic categories and MRS relations
2. Repeat 1. for all words in target language translation
3. Map template mappings onto source and target syntactic categories
 - Translations that match indicate existence of compatible rule template
4. Create a transfer rule by calling the rule template with lists of MRS relations as its arguments

The results of open category transfer rule acquisition from EDICT are summarized in Table 1. We have also extracted several thousand rules of multiple words in length, but they have not been tested yet.

| Rule type | BTEC* vocabulary | Total rules |
|---------------------|------------------|-------------|
| Adj→Verb | 98 | 250 |
| Verb→Adj | 239 | 268 |
| Adj+Noun→Adj+Noun | 478 | 527 |
| Intransitive Verb | 1,273 | 2,519 |
| Noun→Adj.+Noun | 2,262 | 2,787 |
| Adj, Adverb | 2,660 | 3,023 |
| Noun+Noun→Noun | 2,945 | 3,135 |
| Noun→Noun+Noun | 2,100 | 3,588 |
| Noun+Noun→Adj+Noun | 3,974 | 4,482 |
| Transitive Verb | 3,299 | 5,344 |
| Noun+Noun→Noun+Noun | 5,303 | 7,909 |
| Noun | 14,489 | 16,242 |
| Total | 39,120 | 50,074 |

Table 1: Results of automatic transfer rule acquisition from EDICT

| IWSLT 2006 Training data results | | | | |
|----------------------------------|--------------|---|--------|---------------|
| Parsing | 28,175 | / | 42,699 | 65.98% |
| Transfer | 9,355 | / | 28,175 | 33.20% |
| Generation | 5,523 | / | 9,355 | 59.04% |
| Overall | 5,523 | / | 42,699 | 12.93% |

Table 2: Coverage for Jaen on the IWSLT 2006 training data

7.3 Evaluation

We evaluated our system on the development set of the IWSLT 2006 evaluation campaign using the rules we acquired as outlined in this section. Evaluation results are summarized in Table 2. We split all translation pairs into individual sentences on our own, yielding a slightly different number of translation sentences than reported in IWSLT 2006’s data. Currently, we have increased our system’s coverage from a starting point of 1.30% up to 12.93%. In doing so, we are able to translate a large number of sentences with interesting phenomena. Our system’s bottleneck is semantic transfer which succeeds 33.20% of the time in comparison to the over 65% success rate of parsing and near 60% of generation. We report a Bleu score of 0.25 for the sentences we translate over the entire training set using a 3-gram Bleu score and one reference translation.

8 Conclusion

In this paper we showed a semi-automatic approach (automatic generation followed by manual checking) to improve disjunctive entries in the JMdict/EDICT lexicon, and outlined some of the future plans. Because the lexicon is freely available, any improvements will be multiplied by the number of projects that use the lexicon, making even small improvements valuable.

Acknowledgments

This paper is largely based on *Semi-automatic Refinement of the JMdict/EDICT Japanese-English Dictionary* (Bond and Breen, 2007), the section on building transfer rules comes from Nichols et al. (2007).

References

- Ulrich Apel. 2002. WaDokuJT - A Japanese-German Dictionary Database. In *Proceedings of Papillon 2002 Workshop (CDROM)*. (Lexicon at: www.wadoku.de/).
- Francis Bond and James Breen. 2007. Semi-automatic refinement of the JMdict/EDICT Japanese-English dictionary. In *13th Annual Meeting of The Association for Natural Language Processing*, pages 364–367. Kyoto.
- Francis Bond, Stephan Oepen, Melanie Siegel, Ann Copestake, and Dan Flickinger. 2005. Open source machine translation with DELPH-IN. In *Open-Source Machine Translation: Workshop at MT Summit X*, pages 15–22. Phuket.

- Francis Bond and Kentaro Ogura. 2007. Combining linguistic resources to create a machine-tractable Japanese-Malay dictionary. *Language Resources and Evaluation*. (Special issue on Asian language technology).
- J. W. Breen. 2004. JMDict: a Japanese-multilingual dictionary. In *Coling 2004 Workshop on Multilingual Linguistic Resources*, pages 71–78. Geneva.
- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A. Sag. 2005. Minimal Recursion Semantics. An introduction. *Research on Language and Computation*, 3(4):281–332.
- EDR. 1990. Concept dictionary. Technical report, Japan Electronic Dictionary Research Institute, Ltd.
- Christine Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Dan Flickinger. 2000. On building a more efficient grammar by exploiting types. *Natural Language Engineering*, 6(1):15–28. (Special Issue on Efficient Processing with HPSG).
- Dan Flickinger, Jan Tore Lønning, Helge Dyvik, Stephan Oepen, and Francis Bond. 2005. SEM-I rational MT: Enriching deep grammars with a semantic interface for scalable machine translation. In *Machine Translation Summit X*, pages 165–172. Phuket.
- Sanae Fujita and Francis Bond. 2006. A method of creating new valency entries. *Machine Translation*. (to appear).
- Jan Tore Lønning, Stephan Oepen, Dorothee Beermann, Lars Hellan, John Carroll, Helge Dyvik, Dan Flickinger, Janne Bondi Johannessen, Paul Meurer, Torbjørn Nordgård, Victoria Rosén, and Erik Velldal. 2004. LOGON. A Norwegian MT effort. In *Proceedings of the Workshop in Recent Advances in Scandinavian Machine Translation*. Uppsala, Sweden.
- Catherine Macleod, Ralph Grishman, Adam Meyers, Leslie Barrett, and Ruth Reeves. 1998. NOMLEX: A lexicon of nominalizations. In *Proceedings of EURALEX'98*. Liege, Belgium.
- Eric Nichols, Francis Bond, Darren Scott Appling, and Yuji Matsumoto. 2007. Combining resources for open source machine translation. In *The 11th International Conference on Theoretical and Methodological Issues in Machine Translation (TMI-07)*, pages 134–142.
- Lars Nygård, Jan Tore Lønning, Torbjørn Nordgård, and Stephan Oepen. 2006. Using a bi-lingual dictionary in lexical transfer. In *EAMT-2006*, pages 233–238. Oslo.

- Kumiko Ohmori and Masanobu Higashida. 1999. Extracting bilingual collocations from non-aligned parallel corpora. In *Eighth International Conference on Theoretical and Methodological Issues in Machine Translation: TMI-99*, pages 88–97. Chester, UK.
- Kyonghee Paik, Francis Bond, and Satoshi Shirai. 2001. Using multiple pivots to align Korean and Japanese lexical resources. In *Workshop on Language Resources in Asia, NLPRS-2001*, pages 63–70. Tokyo.
- Melanie Siegel. 2000. HPSG analysis of Japanese. In Wolfgang Wahlster, editor, *Verbmobil: Foundations of Speech-to-Speech Translation*, pages 265–280. Springer, Berlin, Germany.
- Jonas Sjöbergh. 2005. Creating a free digital Japanese-Swedish lexicon. In *Proceedings of PACLING 2005*, pages 296–300. Tokyo, Japan. URL `\url{http://www.nada.kth.se/~jsh/publications/jlex.pdf}`.
- Takaaki Tanaka and Yoshihiro Matsuo. 1999. Extraction of translation equivalents from non-parallel corpora. In *Eighth International Conference on Theoretical and Methodological Issues in Machine Translation: TMI-99*, pages 109–119. Chester, UK.
- Masao Utiyama and Mayumi Takahashi. 2003. English-Japanese translation alignment data. <http://www2.nict.go.jp/x/x161/members/mutiyama/align/index.html>.
- Yujie Zhang, Qing Ma, and Hitoshi Isahara. 2005. Automatic construction of a Japanese-Chinese translation dictionary using english as an intermediary. *Journal of Natural Language Processing*, 12(2):63–85. (in Japanese).