Practical and Efficient Organization of a Large Valency Dictionary*

Francis Bond and Satoshi Shirai
NTT Communication Science Laboratories
1-1 Hikari-no-oka, Yokosuka-shi, Kanagawa-ken, Japan 239
bond@cslab.kecl.ntt.co.jp

Abstract

This paper describes the design and ongoing construction of a large bilingual valency dictionary. The first half describes the existing dictionary of 16,000 Japanese-English patterns, how it was built, how it is used, and points out some shortcomings. The second half introduces three proposals, originally put forward by Somers (1987) to improve the dictionary by separating the complement/adjunct distinction from the use of case-roles, organizing the case roles in a grid, and making the English and Japanese Lexicons separate entities, linked by informative links.

1 Introduction

NTT's semantic valency dictionary was built as part of research into Japanese-to-English Machine Translation (Ikehara *et al.* 1991), and a subset of the information (not including case-role and English syntactic information) has been published as Ikehara *et al.* (1997b).

This paper is divided into two sections. First, a description of NTT's semantic valency dictionary, how it was built and some of its uses. Second, are three proposals for extending and improving the dictionary by changing its structure.

2 ALT-J/E's Valency Dictionary

In this section we describe the design and construction of NTT's semantic valency dictionary.

2.1 Description

The valency dictionary is used to describe predicates (verbs, adjectives and copula expressions) both for parsing Japanese and for transferring them into English.

The dictionary entries, which we call patterns, each consist of a predicate, one or more case slots and modal information.

In the Japanese side, case slots are marked with syntactic form semantic constraints.

The syntactic form is given as a phrasal type: clause, noun phrase or adverbial. Only one clause types is allowed: quoted clauses. Noun phrases are listed with the possible particles[1] they may appear with. Adverbs can be time, quantity or other.

The semantic constraints can either be nodes in the semantic hierarchy (Ikehara *et al.* 1997a), or strings that match the surface form. The lower the constraint is in the semantic hierarchy, the better the match. Slots can be explicitly filled, giving idiomatic patterns, which are stored in a separate dictionary.

Because the dictionary was designed for Japanese-to-English machine translation, there are no semantic constraints given for the English side, instead the constraints can be thought of as constraints on the pair of patterns. The English side gives the syntactic form of the translation of the predicate and each case element. The linear order of the English elements, grammatical case, and preposition is also given where applicable. There can be case-slots that have entries only in one language.

An example entry is given in Figure 1. Each entry consists of the Japanese predicate, its English translation and the case-slots. Modal information and some of the detailed syntactic constraints are omitted for brevity. Each slot is labeled with a slot number (S_i), followed by the case-markers, case-role, English surface form and semantic constraints (on a separate line).

The links between the two languages are marked with case-roles (aka deep cases or theta roles),

*This paper was presented at the workshop on Multilingual Information Processing as part of the 4th Natural Language Processing Pacific Rim Symposium 1997: NLPRS-97; Phuket, Thailand

[1]Particles are a closed class of post-positional case markers that mark Japanese noun phrases.

Japanese	case	English
iku$_1$		**go**$_1$
S_1 ga `agent, vehicle, animal`	←N1→	NP Subj
S_2 ni e made `-road, -rail, theatre, places, place`	←N3→	PP *to* Acc
S_3 kara yori `-road, -rail, places, place`	←N4→	PP *from* Acc
iku$_2$		**go**$_2$
S_1 ga `agent, vehicle, animal`	←N1→	NP Subj
S_2 o `places, place`	←N8→	PP *along/around* Acc

Figure 1: Entry for *iku* "go"

which we list in Figure 2.

There are 14 cases and 3 adverbials. The cases can also be used for adjuncts, along with another 10 other more specific cases, of which we list only the three time cases (TN1, TN2 and TN3).

2.2 Construction and Maintenance

The construction and maintenance of the valency dictionary is described in this section. More detail can be found in Shirai *et al.* (1996c).

Creating the dictionary consists of five steps:

Step-1 Identifying needed patterns

Step-2 Constructing translation examples

Step-3 Making the pattern entries

- Parsing the examples and linking the constituents
- Choosing which constituents should be entered
- Adding semantic constraints and case-roles
- Ordering different patterns

Step-4 Checking the pattern produces the correct translation

Step-5 Checking the interaction with other patterns

In the following sections we discuss how the five steps were handled during different phases of the dictionary's construction.

2.2.1 Phase 1

The first 5,000 or so patterns were made by hand. The patterns were constructed by consulting Japanese-to-English dictionaries. In addition, patterns were added as needed whenever there was a problem in translating a sentence.

As the number of patterns grew, it became harder to test them. An **input support system** was built that allowed analysts to check the format of the new entries as they were being built, and then run them through the translation system. This brought the construction time down to 40 minutes per pattern. Using this tool, the system was extended to around 10,000 patterns.

2.2.2 Phase 2

With around 10,000 patterns, we needed to refine steps 1 and 3. Dictionaries for human readers rarely had examples of all the patterns needed for an NLP system so it was hard to identify new patterns. This was particularly a problem for native Japanese verbs, which tend to be more polysemous than Sino-Japanese ones. We therefore decided to systematically go through the Japanese Information-Technology Promotion Agency's set of (IPA 1987).

In addition, a **construction support system** was built, which automatically made a candidate entry by parsing the Japanese example sentence and a specially written English equivalent (of the form 'X goes to Y'). This reduced the time required to make each pattern to around 12 minutes per pattern.

In this phase, there was considerable consolidation of existing patterns, so, while the total number of patterns only increased to 11,000, we estimate that the cover was increased by more than 10%.

2.2.3 Phase 3

The IPAL basic verb list still did not have enough patterns, so a new approach was taken, where the analysts made as many example sentences as they could for each verb, which were then professionally translated, as described in Shirai *et al.* (1996b). This gave many new patterns, and we believe has brought us close to the practical limit of creating new patterns by introspection.

The **construction support system** was extended so that English patterns could be automatically created from raw English text, using the skeleton-flesh approach of Yokoo *et al.* (1994). In this approach, the most common syntactic structures are prepared as skeletons, which are then fleshed out by adding semantic constraints, case-roles and other information such as prepositions.

Label	Name	Particles	Preposition
N1	Agent	ga (kara, towa) [ha]	Subject
N2	Object-1	o (nituite) [ga]	Object
N3	Object-2	ni (...)	Indirect-Obj
N4	Source	kara, yori	from
N5	Goal	ni, e, made	to (until)
N6	Purpose	ni	for
N7	Result	ni, to	as
N8	Locative	ni, o, de, e, kara	in/at/on
N9	Reciprocal	to	with
N10	Quotative	to	
N11	Material	kara, yori, de	with, from
N12	Cause	kara, yori, de	for
N13	Instrument	de	with
N14	Means	de	by
QUANT	Quantity		
TIME	Time		
ADV	Adverb		
TN1	Time-position	ni	at/in/on
TN2	Time-source	kara	since/from
TN3	Time-goal	made	until

Figure 2: Case-roles

In addition, candidate semantic constraints are proposed from the parse of the Japanese sentence. This reduced the time required to produce each pattern down to around 6 minutes per pattern, nearly 7 times faster than the original method.

Using these example sentences and tools, the dictionary was extended to around 16,000 patterns.

2.2.4 Phase 4

Currently, two areas of extension are being explored. The first is construction of domain specific valency dictionaries (Shirai *et al.* 1996a). The second is automatic construction by both the extraction of candidates (Takahashi *et al.* 1997; Haruno and Yamazaki 1996), and the induction of semantic constraints (Akiba *et al.* 1995). It is estimated that we need at least 25,000 patterns to cover around 80% of Japanese verbs (Shirai *et al.* 1995).

2.3 Use

The main use of the valency dictionary is to select the correct dependency structure of the Japanese input, and then to transfer it to an English structure, as described in Ogura *et al.* (1993).

To determine the dependency structure, Input sentences are first analyzed by a morphological analyzer, which separates the text into words, marked with part of speech and multiple senses. The output of this is then parsed to give candidate dependency structures. These are then matched with the idiomatic pattern dictionary, and then the general valency dictionary.

When there are multiple candidates for the predicate, the matches are weighted using the following criteria: does an input word match an explicit entry in the dictionary (e.g. an idiom). If not, choose the pattern with the highest total score. Each matching slot is given a value according to the level of the matching semantic constraint (from 100 at a leaf level to 60 for the top level). This is then adjusted according to the case-role: N2 and N3 are increased the most, N1 is increased a little, N4 and N5 are unchanged, N6–N8 are decreased slightly and the rest are decreased even more. This reflects the strength of the case-element's connection to the predicate. The pattern's total score is then the sum of the scores of its elements.

Once the highest ranking pattern has been chosen, it gives the backbone of the dependency structure. The constraints given by the predicate to the case-elements are then used to disambiguate the case-elements themselves.

The main use of the case-roles is to link the Japanese and English patterns. They are also used to select prepositions for the adjunct cases, although the default prepositions can be overwritten for adjuncts entered in the dictionary.

The case-roles (slightly augmented) are also

used to determine the order of English adverbs (Ogura *et al.* 1997:p 22). Complement elements come closer to the predicate than adjuncts, and adjuncts are ordered as follows:

> Manner < Means < Instrument < Position < Direction < Time-position < Time-duration < Frequency

Another use of the valency dictionary is in the generation of articles. Temporal case elements are generated with special rules (Bond *et al.* 1997), and locative case elements are definite by default.

The case-roles thus serve as useful links between the two languages, as well as serving as triggers for some general rules.

There are however, some problems with the case-roles, which will be discussed in the next section.

2.4 Some Problems

In practice, the major problem with using case-roles is that it is hard for analysts to assign values to entries in the lexicon. Many natural language systems use case-roles, but there is little agreement as to how many there should be, let alone what they should be. A good example of this problem is the LUTE system, which ended up with incompatible sets of cases for Japanese (29 case-roles in 6 groups) and English (42 case-roles). Another set is defined by Nomura and Muraki (1996:p 645) with 34 case-roles (deep cases) with 16 used in the dictionary.

The choice of 24 case roles for ALT-J/E, with only 14 used in the dictionary was a pragmatic one, this was the number of cases that seemed necessary, and that could easily be distinguished by the analysts. As can be seen in Figure 1, the choice is sometimes questionable, slot-2 of iku_2 should be N5 (Goal) rather than N3 (Object-2). N3 is however the default for the case-marker *ni*, and was assigned instead.

There are also problems caused by the conflation of the degree of valency (how closely related the case-element is to the verb) and the case-role. The accusative-case (marked by *o*) in a verb such as *tazuneru* "visit" is obligatory, so should be marked with N2, but should be locative, which calls for N8. Due to this conflict in the definition, some verbs of this type are marked as N2, and some N8. This is a problem because such arguments should be definite by default (a rule which is triggered by N8), but allow floating quantification (a rule which is triggered by N2).

Another problem, in practice, is the close association of case-roles with their surface markers. Some rules are written using N1 and N2 to mean nominative and accusative surface case. Ideally during processing at least three levels of information about case are needed: the surface case marker or markers, the canonical case marker (from the lexicon) and the case-role.

Finally, the direct linking of the two languages means that any differences in predicate meaning in English, have to be anticipated during the Japanese processing. There is no chance to delay the choice of English predicate, and the Japanese parse can be quite counter intuitive.

3 A Different Approach

In this speculative section, I make three proposals to improve the structure, and ultimately simplify the maintenance and construction of the dictionary. The proposals are to: separate the degree of valency from the case role (§ 3.1), recast the case-roles as a case-grid (§ 3.2) and treat the Japanese and English lexicons as separate entries, with informative links between them. All three proposals were originally proposed by Somers (1987), although we modify his proposals somewhat, and offer more justification for them, based on our own experiences.

Each proposal could be implemented separately, but they all fit together.

3.1 Separate Valency from Case

The first proposal is to add a new variable for each slot: the degree of variable binding, which shows how closely an element is connected to the predicate of the clause it appears in. Somers (1987:p 266) proposed a 6 valued variable; we propose adding another value for Pustejovsky's (1995:pp 63–67) shadow arguments, bringing the number to 7. Note that the degree of variable binding is used for items marked in the lexicon (marked in bold) as well as those determined during parsing, such as adjuncts.

0 Zero complement (contained in predicate)

1 Integral complement (idiom dictionary)

1.5 Shadow complement

2 Complement

3 Middle

4 Adjunct

5 Extra-peripheral

Integral complements are obligatory parts of idioms like *the bucket* in *kick the bucket*. They cannot be removed without changing the meaning of

the verb. Shadow complements are elements that are only expressed if they are special in some way, such as *with butter* in the verb *butter*. It is strange to say *butter the bread with butter*, but *butter the bread with expensive butter* acceptable.

Complements are the normal obligatory arguments of the verb, such as the subject of the verb *go*. Middles are elements that are strongly associated with a verb, but not obligatory, such as *to school* in *go to school*, or *with a hammer* in *break a glass with a hammer*. An independent argument for the existence of elements such as these, between true complements and adjuncts is given by Verspoor (1996) in a treatment of the semantic contribution of prepositional phrases, who refers to them as pseudo complements. The addition of a middle value makes the job of dictionary analysts much simpler, particularly for Japanese, where free omission of most elements makes it hard to decide whether they are obligatory or not.

Adjuncts are optional sentence elements, corresponding to Ogura *et al.*'s (1997) adjuncts. Extra-peripheral elements are sentence modifying elements, such as Ogura *et al.*'s (1997) disjuncts and conjuncts. We give some examples in the next sentences:

(1) Fred$_2$ kicked [the bucket]$_1$

(2) Unexpectedly$_5$, [the bread]$_2$ was buttered [with expensive butter]$_{1.5}$ [this morning]$_4$

(3) He$_2$ did not go [to work]$_3$ [with you]$_4$

As a first approximation, values in the current dictionary would be mapped as follows: explicit entries (in the idiom dictionary) map to integral complements; N1-N3 map to Complements, the rest map to Middles. Complements should be weighted as higher than Middles, perhaps 1.5 to 1. Other weightings can also be made according to either case-role, or Japanese surface particle, as required.

Somers (1987:p 267) speculates that when mapping from one language to another, elements would either map onto elements with the same degree of valency, or one more or less. Thus a middle maps onto a middle, complement or adjunct and so on. For that reason, a zero complement is proposed to account for matches such as English *take part* and French *participer*; *part* matches to an empty zero complement in *participer*. There are however, many examples between Japanese and English where the mapping is across degrees that differ by more than one, such as the combination of a Japanese verb and manner adverbial translated as an English verb: for example *burabura aruku* "relaxed walk" with *stroll* (degree 4 mapping to degree 0). Therefore we need to allow links between elements with any degree of valency.

Note that our lexicon could also be extended by adding a wider variety of clause types as complements, such as adjective phrases, but that would be a separate exercise.

The different degrees of valency can be used to explain sentences where the same case role appears twice, breaking the one-case-per-argument condition of Fillmore (1968:p 24). Consider the following examples, from Somers (1987:p 192):

(4) Taroo-no otoosan-ga shinde-shimatta
Taroo-ADN *Father*-NOM *dying do past*
Taroo's father died

(5) Taroo-ga otoosan-ga shinde-shimatta
Taroo-NOM *Father*-NOM *dying do past*
Taroo's father died on him

In the second sentence, Taroo can be thought of as the subject of the whole sentence, and should have a higher degree of valency than *otoosan* "father". As *otoosan* "father" is a normal complement of *shinu* "die" its degree of valency is 3. The external subject Taroo is a peripheral adjunct, and would have a valency of 6. This sentence could also possibly be explained as different case-roles: Taroo as experiencer and *otoosan* "father" as agent.

3.2 Organize Cases in a Grid

The second proposal, again closely following Somers (1987:chapter 10), is to regularise the case roles in a grid, given in Figure 3. The grid has two major attractions. First, it puts a well defined limit on the number of possible cases. Second, it allows generalizations to be made along columns and rows. Both of these make it easier to assign case roles.

The columns consist of the four localist values, exemplified by the locative row. The four values are the source, the path taken, the goal, and a point (possibly along the path. The rows are more guided by semantic criteria. Very broadly, the Active row represents actions, where as the Objective row represents processes. The dative row is used for psychological and possessive predications.The temporal and locative rows are self explanatory. Finally, the Ambient row is for more abstract cases such as reason, manner and aim. In Figure 3, each cell has been given a short descriptive name as aguide (two names for the dative row). More detailed descriptions are given in Somers (1987:pp 200–206).

	Source	Path	Goal	Local
Active	Instigator N1	Means N14	Recipient N6	Patient N9
Objective	Material N11	Instrument N13	Result N7	Changed N2
Dative	Stimulus (N1)	Medium (N2)	Experiencer N3	Content N10
	Owner	Price	Recipient	Transferred
Locative	Source N4	Path (N8)	Goal N5	Point N8
Temporal	From TN2	Duration (TN1)	Until TN3	When TN1
Ambient	Reason N12	Manner (N14)	Aim (N7)	Condition (N11)
Particle	**ga**, wa, towa kara, yori	**de**, to, o	**ni** e made to	**o**, to ni, de, nituite
Preposition	**from** for	**by**, with for, around	**to**, until as, for	**in/at/on** with, about

Figure 3: Case Grid

As a rough guide to assigning cases to the grid, the column can be assigned on the basis of the preposition or case-particle, although obviously it is not a one-to-one mapping. The row has to be determined by the verb and case-element meaning. The same table can be used for complements (for both nouns and predicates) as well as adjuncts, although some extra ones may be needed for adjuncts.

The grid allows easy identification of locative and temporal expressions for the generation of articles and prepositions.

It would be worth experimenting by weighting the columns and rows differently for the stren1gth of the match, for example Objective and Dative increased, Temporal and Locative decreased, and maybe Source and Goal increased slightly. Of course, actual values for weights need to be obtained empirically.

As an example of the ease of use of the proposal, we apply it to a difficult class of verbs, including verbs of potential: *wakaru* "understand", perception *kikoeru* "can hear" and desire: *hituyō-da* "need", which have the subject marked with *ni* (which is normally dative or locative) and the role equivalent to the English object marked with *ga*, the nominative case. The case grid allows us to mark these roles reasonably intuitively. The subject is marked as Dative-Goal (or experiencer), where the Goal matches the *ni* marking, and the other role is marked as the dative source (or stimulus), the thing that stimulates the senses or motivates desire.

(6)

predicate:			wakaru	
S_1	NP	3	ni∧ga	Dative-Goal
S_2	NP	3	ga	Dative-Source

(7)

predicate:			au	
S_1	NP	3	ga	Active-Source
S_2	NP	3	ni	Active-Local
S_2	NP	3	to	Objective-Local
∧				
S_1	NP `plural`	3	ga	Active-Source

For these verbs the subject can also be marked with *ga* probably because of the normal association of the nominative case with the subject. We therefore mark the subject as taking either *ni* or *ga* with the same case-role.

Ideally case-roles are the same for equivalent verbs across languages, and this should be the default in links. However, it is not always the case, particularly for languages as different as Japanese and English, so there must be a way of linking a slot marked with one case-role in one language to one marked with a different one in the other.

The case-roles can be thought of as upper nodes of a hierarchy of more detailed semantic roles such as GIVEE, DONATEE, HANDEE (all subsumed by Dative-Goal), as proposed by Pollard and Sag (1994:pp 342-343). They are useful in two ways: one, to make generalizations over classes; two, as a first level of information until more detailed descriptions become available.

3.3 Separate Languages

The final proposal is to store English and Japanese patterns separately. The valency dictionary has been constructed this way from the beginning, but it has not been exploited fully.

One advantage of storing the English and Japanese language patterns separately is that the separate lexicons can be learned from monolingual corpora, which are always larger and more easily available than bilingual corpora. Another advantage is that it should become easier to to make

generalizations within each language, as well as easier to eliminate inconsistencies.

A potential disadvantage is that it reduces the savings that can be made by simplifying the target language dictionary in a one way transfer system. As generation becomes more sophisticated, however, more information is wanted in the target dictionary anyway; and for a two way system, this information would be necessary in the first place.

The links between the two dictionaries have to be informative, not just matching one predicate with another. By judicious use of defaults, the links can be quite small: by default, the same case-roles should link to each other $\pm$ 1 up and down the valency binding. There has to be a provision to explicitly link slots, and even add constraints on the linked slots, to handle mismatches between the languages. The links can be thought of as bilexical rules, along the lines of those proposed by Trujillo (1995).

If the two monolingual dictionaries exist already, then linking them can start off as a simple progress of linking predicates that match in a bilingual corpus, and analysts only have to examine those that don't fit the default match parameters. By considering links along rows and columns, candidate links can be suggested even for those that do not match well, and of course much more could be done if the bilingual information was richer.

Similarly, if a monolingual dictionary with case-roles and some bilingual data exist, they can be used to boot strap a dictionary in another language.

The combination of two monolingual dictionaries and links can either be precompiled into a single transfer dictionary, or treated as three separate entities. The advantage of precompilation is that it gives a dictionary equivalent to the existing one, so that the changeover could be made seamlessly. In the compiled dictionary, the constraints on each element would be a combination of the strictest from the two patterns and link.

More advantages could be gained by keeping the dictionaries separate. In this case the target predicate does not have to be chosen until later in the transfer/generation stage making it easier to apply purely target language constraints such as collocational constraints.

Finally, the process of construction may become simpler, and will definitely be more consistent if two monolingual dictionaries and links are used. Assuming that the pattern to be entered has been identified, step-3 of the construction process can proceed as follows:

1. If there is no Japanese pattern, make one (possibly automatically)

2. If the candidate translation is known, create potential links for analyst to select.

3. If the translation is unknown, suggest candidates with similar case-roles and valency bindings, ordered from exact match down.

4. If no suitable English pattern exists, make one (possibly automatically), or modify an existing one

5. If a new English pattern has been made, suggest candidate translations from the Japanese patterns

This procedure takes advantage of existing knowledge, if a suitable entry exists, in either language, then it can be used directly, and just linked. The links add information for disambiguation, but only when it is needed. For example, *warau* "laugh/smile" does not need to be disambiguated in Japanese analysis, only in translation. In addition, as potential links are checked for new patterns in both directions, the coverage should be better than only checking one way. Verb sense hierarchies, such as that proposed by (Nakaiwa *et al.* 1994), could be used to constrain possible candidates for linking.

4 Conclusion

We have a large and very useful bilingual valency dictionary. It was hard to build, is sometimes inconsistent, and is quite hard to extend. To improve the quality of translation, we need to extend it in at least two ways: size and complexity of information. We propose that we can make it both more useful, and easier to build, by extending case-roles, separating case from valency binding, and treating different languages separately in the lexicon.

Acknowledgments

The authors would like to thank the designers and maintainers of the pattern dictionary. Much of the work of designing the dictionary was done by Akio Yokoo and Satoru Ikehara. The actual construction of patterns has been done by analysts including Satsuki Abe, Hiroko Inoue and Izumi Watanabe. We would also like to thank Tim Baldwin, Osamu Furuse, Yoshihiro Matsuo, Kyonghee Paik and Kentaro Ogura for their comments and discussion.

References

Akiba, Yasuhiro, Megumi Ishii, Hussein Almuallim, and Shigeo Kaneda. 1995. Learning English verb selection rules from hand-made rules and translation examples. In *Sixth International Conference on Theoretical and Methodological Issues in Machine Translation: TMI-95*, 206–220, Leuven.

Bond, Francis, Kentaro Ogura, and Hajime Uchino. 1997. Temporal expressions in Japanese-to-English machine translation. In *Seventh International Conference on Theoretical and Methodological Issues in Machine Translation: TMI-97*, 55–62, Santa-Fe.

Fillmore, Charles J. 1968. The case for case. In *Universals in Linguistic Theory*, ed. by Emmon Bach and Robert T. Harms, 1–88. New York: Holt, Rinehart and Wilson, Inc.

Haruno, Masahiko, and Takefumi Yamazaki. 1996. High-performance bilingual text alignment using statistical and dictionary information. In *34th Annual Conference of the Association for Computational Linguistics*, 131–138.

Ikehara, Satoru, Masahiro Miyazaki, Satoshi Shirai, Akio Yokoo, Hiromi Nakaiwa, Kentaro Ogura, Yoshifumi Ooyama, and Yoshihiko Hayashi. 1997a. *The Semantic System*, volume 1 of *Goi-Taikei — A Japanese Lexicon*. Tokyo: Iwanami Shoten.

——, ——, ——, ——, ——, ——, ——, and ——. 1997b. *The Valency Dictionary*, volume 5 of *Goi-Taikei — A Japanese Lexicon*. Tokyo: Iwanami Shoten.

Ikehara, Satoru, Satoshi Shirai, Akio Yokoo, and 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*, 101–106, Washington DC. (cmp-lg/9510008).

IPA. 1987. IPAL (basic verbs). Lexicon, Information-Technology Promotion Agency, Tokyo, Japan. (ftp://ftp.mgt.ipa.go.jp/pub/ipal).

Nakaiwa, Hiromi, Akio Yokoo, and Satoru Ikehara. 1994. A system of verbal semantic attributes focused on the syntactic correspondence between Japanese and English. In *15th International Conference on Computational Linguistics: COLING-94*, 672–678, Kyoto.

Nomura, Naoyuki, and Kazunori Muraki. 1996. An empirical architecture for verb subcategorization frame. In *16th International Conference on Computational Linguistics: COLING-96*, 640–645.

Ogura, Kentaro, Francis Bond, and Satoru Ikehara. 1997. A method of ordering English adverbials — as exemplified in Japanese-to-English machine translation —. *Journal of Natural Language Processing* 4.17–39.

Ogura, Kentaro, Akio Yokoo, Satoshi Shirai, and Satoru Ikehara. 1993. Japanese to English machine translation and dictionaries. In *44th Congress of the International Astronautical Federation*, Graz, Austria.

Pollard, Carl, and Ivan A. Sag. 1994. *Head Driven Phrase Structure Grammar*. Chicago: University of Chicago Press.

Pustejovsky, James. 1995. *The Generative Lexicon*. MIT Press.

Shirai, Satoshi, Satoru Ikehara, Akio Yokoo, and Hiroko Inoue. 1995. The quantity of valency pattern pairs required for Japanese to English machine translation and their compilation. In *Natural Language Processing Pacific Rim Symposium '95: NLPRS-95*, 443–448.

Shirai, Satoshi, Hiroko Inoue, Noriko Itakura, Satoru Ikehara, and Akio Yokoo. 1996a. A semantic valency dictionary for technical domains. In *2nd Annual Meeting of the Association for Natural Language Processing*, 13–16. (in Japanese).

Shirai, Satoshi, Hiroko Inoue, Ode???, Noriko Itakura, and Akio Yokoo. 1996b. Compiling valency pattern pairs for Japanese-to-English machine translation based on the IPAL basic verb examples. In *53rd Annual Convention of the IPSJ*, volume 4L-4, 2:59–60. (in Japanese).

Shirai, Satoshi, Hiromi Ueda, Tomiko Hyodo, Akio Yokoo, and Satoru Ikehara. 1996c. Computer aided design of valency pattern pairs for Japanese-to-English machine translation. In *IEICE Technical Report NLC96-34*, 25–30. IEICE. (in Japanese).

Somers, Harold. 1987. *Valency and Case in Computational Linguistics*. Edinburgh University Press.

Takahashi, Yamato, Satoshi Shirai, and Francis Bond. 1997. A method of automatically aligning Japanese and English newspaper articles. In *Natural Language Processing Pacific Rim Symposium '97: NLPRS-97*, 657–660.

Trujillo, Arturo. 1995. Bi-lexical rules for multi-lexeme translation in lexicalist MT. In *Sixth International Conference on Theoretical and Methodological Issues in Machine Translation: TMI-95*, 48–66.

Verspoor, Cornelia Maria. 1996. A perspective on PPs. In *Edinburgh Working Papers in Cognitive Science, Vol. 12: Studies in HPSG*, ed. by Claire Grover and Enric Vallduví, volume 12, chapter 7, 229–271. Centre for Cognitive Science, University of Edinburgh.

Yokoo, Akio, Hiromi Nakaiwa, Satoshi SHirai, and Satoru Ikehara. 1994. Skeleton-flesh type semantic structure dictionaries for Japanese-to-English machine translation. In *48th Annual Convention of the IPSJ*, volume 6Q-8, 3:139–140. (In Japanese).