

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

An Overview of NLP in ALICE-chan

D.A. Evans, L.S. Levin, M. Thurn, S.K. Handerson,
K. Horiguchi, K.C. Scarpinato, D. Varma

December 1993

Report No. CMU-LCL-93-3

Laboratory for Computational Linguistics

139 Baker Hall
Department of Philosophy
Carnegie Mellon University
Pittsburgh, PA 15213

Contents

1	Character-Independent Processing	1
2	The NLP Lexicon	3
3	Morphological Analysis	6
4	Segmentation	7
5	Analysis of Syntactic Structure	8
6	The Parser	9
7	The Grammar	9
8	The Mapper	10
9	The Disambiguator	11
10	The Matcher	12
11	Overview of Error Detection and Feedback	13

List of Figures

1	An example of Japanese text entry.	2
2	The NLP lexicon entry for the Japanese word “hon” (“book”).	3
3	The NLP lexicon entry for the Japanese word for “violin.”	3
4	The NLP lexicon entry for the Japanese proper noun “Yoshiko.”	3
5	The NLP lexicon entry for the stem of the Japanese verb “kaku” (“write”).	4
6	The NLP lexicon entry for the Japanese verb morpheme “ki”.	4
7	The NLP lexicon entry for the Japanese present formal verb morpheme “masu”.	4
8	The NLP lexicon entry for the stem of the Japanese morpheme “i.”	4
9	The NLP lexicon entry for the Japanese verb gerund morpheme “te”.	4
10	The NLP lexicon entry for the Japanese past informal verb morpheme “ta”.	4
11	The NLP lexicon entry for the stem of the Japanese auxiliary verb “iru”.	4
12	A Japanese sentence meaning “Keiko ate sushi.”	8
13	The raw feature-structure resulting from parsing the sentence above.	8
14	A rule from the parser’s grammar.	9
15	Some disambiguation dialogs.	11

⌘	Feature-slot for the author's sentence.	14
17	Feature-slot for the student's sentence.	14
18	The feedback produced after matching the above two feature-slots.	M

List of Tables

I	The Japanese word for "Tokyo" in various character formats.	1
----------	---	---

1 Character-Independent Processing

One of the most design decisions for ALICE was that the system will operate on input in *hiragana*, *kanji*, or *romaji*. The NLP system is designed to accept any mixture of *romaji* and Japanese characters as input, and returns a mixture of English and Japanese in its output. Thus, the student can type her answer in *romaji*, *hiragana*, *kanji*, or any mixture of these. Feedback to the student is presented as English sentences sprinkled with references to the student's response (possibly in Japanese characters).

Voice-inputting Japanese characters, ALICE-chan utilizes the common Japanese input method whereby the user types *romaji* on a standard English keyboard, the *romaji* is converted to *hiragana* automatically as the user types, and the user can request that *hiragana* be converted to *kanji* by pressing a special key. Other functions include scrolling through alternative *kanji*; and stretching or shrinking the length of word segments, which controls how many *kana* syllables will be grouped together for conversion into *kanji*.

Romaji-to-hiragana conversion is carried out by a finite state transition process written in C. *Hiragana-to-kanji* conversion is carried out by external function calls to *jsrvr*, a public domain *kanji*-conversion server.

<i>English</i>	Tokyo	5 characters, 5 bytes
<i>romaji</i>	toukyou	7 characters, 7 bytes
<i>hiragana</i>	とうきょう	5 characters, 10 bytes
<i>kanji</i>	東京	2 characters, 4 bytes

Table 1: The Japanese word for "Tokyo" in various character formats.

Options Input-Mode

《きょう》

漢

The user has typed "kyous".

Options Input-Mode

《きょうするしごとはありません》

漢

The user has typed "kyousurushigotohaarimasen".

Options Input-Mode

《 供する 仕事は ありません 》

上

The user has pressed the spacebar and jserver has given its first guess at segmentation and kanji selection.

Options Input-Mode

《 キヨウス ルシ 五戸は ありません 》

上

The user has typed "Ctrl-i" to shrink the length of the first segment by one kana, and jserver has given its first guess at kanji selection for the new segmentation.

Options Input-Mode

《 今日 する 仕事は ありません 》

上

The user has typed "Ctrl-i" again to shrink the length of the first segment by one more kana, and jserver has given its first guess at kanji selection for the new segmentation.

Options Input-Mode

今日する仕事はありません

漢

The user has typed "Return" to solidify the entire sentence of kanji.

Figure 1: An example of Japanese text entry.

2 The NLP Lexicon

The NLP lexicon contains information that allows the system (1) to recognize words in all of their morphological variants and (2) to identify syntactic and semantic features of the word.

Each lexical entry consists of two main parts, one is a list of keys in *romaji*, *kana*, and *kanji*, and the other is a list of syntactic and semantic features. The keys indicate possible orthographic realizations of a word. For example, the word in Figure 2 can appear as *hon* (romaji), *hon*, (hiragana), or *hon* (kanji). If a word can be written in more than one way, all the possible keys are listed as an embedded list. For example, the word that means "violin" (Figure 4) can be written in two ways in *katakana*, therefore there are two *katakana* strings embedded in the key list. Lexical entries for proper names often include many *kanji* variants, as shown in Figure 4.

```
(("hon" "hon" "hon" "hon")
 (S "book" M (hon) COUNT volume CAT noun INANI-SUBJ +
  BL + R (FINAL PART) L (BEGIN G-HON)))
```

Figure 2: The NLP lexicon entry for the Japanese word "hon" ("book").

```
(("baiorin" ("hon" "hon" "hon" "hon")
 (S "violin" M (baiorin) COUNT machine CAT noun INANI-SUBJ +
  BL + R (FINAL PART) L (BEGIN)))
```

Figure 4: The NLP lexicon entry for the Japanese word for "violin." Notice the multiple strings representing the various *katakana* versions of the word.

```
(("yosiko" "yosiko" "yosiko" "yosiko" "yosiko" "yosiko" "yosiko" "yosiko")
 (S "Yoshiko—girl's name" M (yosiko) CAT noun INANI-SUBJ +
  ANIMATE + PROPER + HUMAN + BL + L (BEGIN) R (SAN NAME FINAL PART)))
```

Figure 4: The NLP lexicon entry for the Japanese proper noun "Yoshiko." Notice the multiple strings representing the various *kanji* versions of the word.

The second part of a lexical entry contains a list of feature-value pairs. The feature *S* stands for *sense*. Its value is a short English gloss indicating the meaning of the word. The feature *M* (morpheme) is a *romaji* spelling of the citation form of the morpheme.

Morphological analysis is guided by the special features *L* and *R*, whose values are names of states in a non-deterministic finite state transition network. The value of *L* is a set of possible starting states for a transition and the value of *R* is a set of possible ending states. The keys of the lexical entry are the possible input strings which license the transition. Figures 5 through 7 show lexical entries involved in the analysis of several forms of the verb *kaku* (write). For example, processing of *kakimasu* starts at the initial state (BEGIN), and uses the input *ka* to get to state *K*. From state *K*, the input *ki* allows a transition to state *MAS*. From state *MAS*, the input *masu* allows a transition to state *FINAL*.

The addition of the lexical entries shown in Figures 8 through 11 allow the system to

((^Mka^M "k*" "k")
 (S "write" M (ka) WO-PATIENT + NI-GOAL + COMP + CAT verb BL +
 ROOTKEYS dummy UNUM dummy SUFKEYS ("ku^M "k" ^Mk) L (BEGIN HON) R (K)))

Figure 5: The NLP lexicon entry for the stem of the Japanese verb "kaku" ("write").

((("KI" "k") (M (ki) L (K) R (P-NI MAS ASPV HON DESIDER)))

Figure (3): The NLP lexicon entry for the Japanese verb morpheme ^b'kr.

((("MASU" "if") (M (masu) STYLE formal TENSE present L (MAS) R (FINAL)))

Figure* 7: The NLP lexicon entry for the Japanese present formal verb morpheme ^wmasif\

process *kaiicimasu* ("write/* gerund + auxiliary "to be") and *kaila* ("write/* past informal), which illustrate the changing of the final stem consonant *-k-* to *-/-* before the gerund (*-tr-*) and past tense (*la*) morphemes. Processing of these words also starts in the initial stair, using the input *ka* for a transition to state *k*. Then, the input */* allows a transition to stair T-GER or T-PAST, etc. Notice the treatment of the stem allomorphs *kak-* and *kai-*. The lexical (Mitrirs for *-ki-* and *-/-* both start in state *k*. but they end in different stairs, rrrllrlin<; thr different possible continuations of each stem.

((("I" 'HV<) (M (i) L (K) R (T-GER T-PAST)))

Figure 8: The NLP lexicon entry for the stem of the Japanese gerund morpheme ^v'kr.

((("TE" »X") (M (te) GERUND + ADVERBIAL - L (T-GER) R (FINAL REQUEST WA MO GER-ASP)))

Figure 9: The NLP lexicon entry for the Japanese gerund morpheme ^v'kr.

((("TA^{II}" "it") (M (ta) STYLE informal TENSE past L (T-PAST) R (FINAL TENSE)))

Figure 10: The NLP lexicon entry for the Japanese past informal verb morpheme ^v'kr.

((("I" "v>")
 (M (i) A (doing) GERUND - L (GER-ASP)
 R (V-NEG MAS DESIDER V-CAUS PL-PRES IMP T-PAST T-GER V-PASS POIEN V-VOLIT V-PROV)))

Figure 11: The NLP lexicon entry for the stem of the Japanese auxiliary verb ^v'kr.

Notice also, that lexical entries do not necessarily correspond to whole morphemes. *Ka-*, *-ki-*, and *-/-*. for example, are parts of the *kak-jkai-* morpheme, and *-/-* might even span morpheme boundaries if the analysis of *kakimasu* is *kak+imasu*. Lexical entries actually

correspond more closely to Japanese orthography than to a linguistic morphological analysis. Each entry corresponds to a whole number of syllables which are the pronunciations of *kanji* or *hana* characters. Since a syllable can be part of a morpheme, or can span morpheme boundaries, character boundaries do not correspond to morpheme boundaries. More details about morphological processing are found in the chapter on morphology.

There are currently about 1700 entries in the ALICE-chan NLP Lexicon.

j

3 Morphological Analysis

The morphological processor is conceptually an interpreter for a nondeterministic finite state automaton, specified by the lexicon. This allows the process to be simple, fast, and flexible (lexicon-driven). This simplicity and speed motivated combining morphological processing with segmentation as a single initial pass over the input string. For sentences with successful parses, the parser is limited to determining the grammaticality of a segmentation, instead of also having to determine a sentence's segmentability. In this section, we will refer to "morphemes" and "words" without the double-quotes, a morpheme simply being a string entered in the lexicon and a word the result of a successful morphological analysis.

Morphemes entered in the lexicon describe a nondeterministic finite-state automaton (NFA) where the transitions are the morphemes. Each morpheme also specifies the allowable prefix states and the potential result states. Morphemes may be entered multiply in the lexicon to express different sets of applicable contexts and result contexts. There are special states signifying the initial state (morphemes that can begin a word) and final state (morphemes that can end a word). A word, according to our morphological process, is any sequence of morphemes that cause the NFA to reach a final state.

During analysis, a simple feature-value pair *feature structure* is built from information associated with the morpheme in the lexicon. This information reflects syntactic information and is converted into a form used by the parser. This information can also be used later, by application-specific processes (for instance, to facilitate user feedback). Since the processor (or NFA) lacks the ability to do even simple checks based upon the built-up features and values, this information is not used directly to enforce morphological constraints, but could be so used by the parser.

Currently, the runtime lexicon is stored in a hash table indexed by all morpheme strings and their prefixes, but other implementations are possible (for instance, a trie). All that is necessary is to know whether a given prefix string can be continued to form a morpheme. A hash table implementation was chosen because the lexicon contains *kanji* and *hiragana* as well as *romaji*, and so a space-efficient trie representation would be difficult, but we have not studied this to any great extent.

4 Segmentation

A problem with NLP for Japanese is that the written language does not have spaces separating words as in most Western languages. One of the **first** problems a parser (human or machine) for Japanese faces is to determine a sentence's word boundaries. We have found that the left-to-right nature of our morphological analysis, together with a few assumptions (or observations) concerning the input sentence, yield a "knowledge-poor*" (i.e. simple*) algorithm which is nonetheless very efficient. The advantage of segmenting prior to parsing is that many words can be ruled out from further consideration, since they cannot be part of a complete¹ sentence parse.

Obviously, in the absence of misspellings, any correct parse of the input sentence must divide the input into non-overlapping words. If we first find all possible words that begin at the first position, the second word in the intended interpretation (parse) must end after one of these possible first words, etcetera. Also, the intended interpretation must have its last word end at the end of the sentence: there are no "extra characters" at the end. Note that there is an implicit assumption that the segmenter knows whether a given string can be continued to form a new word: otherwise the "find all possible words" step might have to look at the entire remaining sentence, for every character position. These are the main assumptions we make about the input: all other information is in the morphology and the lexicon.

These constraints translate into an algorithm with a single left-to-right word generation phase and a right-to-left elimination or filtering phase. Note that words need not be generated for every character position: only those "reachable" from the beginning of the sentence. The filtering phase is simply the right-to-left analogy of the generation phase, keeping in mind that we do not need to add words that are not "reachable" from the beginning of the sentence.

All interesting but not crucial last step constructs the input to the parser, not as words or characters, but "chunks". A "chunk*" is a sequence of characters that is never divided between words in any segmentation. Note that passing words to the parser would be difficult or impossible (except in a chart parser), since there may be several segmentations with different word boundaries. Parsing a sequence of "chunks", however, is conceptually similar to parsing a sequence of characters with an LR parser. Also, the information needed to construct the "chunks" is already present in the segmentation data structures.

Unfortunately, this approach does not work well in the case where there are spelling errors. If the spelling error is such that there are correct segmentations, this would make such spelling errors harder to detect. A good approach might be to first assume no spelling errors, and if either there are no segmentations, no parses for the resulting segmentations, or many or serious grammatical or exercise-specific errors, then a more complex and time-consuming process ensues, that starts by trying to identify expected words in the input.

5 Analysis of Syntactic Structure

The goal of syntactic analysis is to identify (1) the predicate of each clause, (2) the predicate's syntactic and morphological features, and (3) a grammatical function and/or semantic role for every other element of the clause. An example of a sentence and the raw form of its analysis are shown in Figures 12 and 13.

Syntactic analysis is used during exercise authoring and during student exercises. In authoring, it is used for analyzing correct answers so that the syntactic analysis of the answers can be annotated by the author, and stored for later comparison to student answers. In student exercises, syntactic analysis is used for analyzing student answers so that they can be compared to previously stored authored answers.

恵子さんが寿司を食べました。

Figure 12: A Japanese sentence meaning "Keiko ate sushi."

```
(( (CAT VERB) (SUFKEYS (*LIST* "ru" "S" "&")) (BL +) (WO-PATIENT +) (S #:EAT)
 (UNUM 3) (ROOTKEYS (*LIST* "tabe" "fz'<" "&'<")) (TENSE PAST)
 (M (*LIST* TABE MASITA)) (STRKEYS (*LIST* "jɛ'<" "&lt;l"))
 (PARTSTRKEYS (*LIST* "&K" "&lt;z")) (ALLKEYS (*LIST* #:G13417 #:G1598))
 (MFSS (*LIST* #:G13418 #:G1599)) (END-CHUNK 3) (BEGIN-CHUNK 3)
 (STR "&'<&lfz") (END-PHRASE NIL) (DICT-FORM "•&'<.ɛ>") (BEGIN-PHRASE 0)
 (SUBCAT
 ((GA ((ACTOR ANIMATE))) (WO ((ACTEE ANY))) (TO ((ACCOMPANIMENT ANIMATE)))
 (DE ((LOCATION PLACE))) (NI ((TIME-AT TIME-NI)))
 (KARA ((ACTOR ANIMATE) (TIME-FROM TIME))) (MADE ((TIME-TO TIME)))
 (WA ((ACTEE ANY) (ACTOR ANIMATE))) (MO ((ACTEE ANY) (ACTOR ANIMATE))))))
 (STYLE FORMAL)
 (ACTOR
 ((BL +) (HUMAN +) (PROPER +) (ANIMATE +) (INANI-SUBJ +)
 (S #:IKEIKO-GIRL'S NAMED (TITLE (*LIST* MR./MRS./MS./MISS)) (RESTRICT +)
 (M (*LIST* KEIKO SAN GA)) (STRKEYS (*LIST* "M?" "S/v" "#"))
 (PARTSTRKEYS (*LIST* "S ɪ̄" "J/u" "-ɛ̄"))
 (ALLKEYS (*LIST* #:G5744 #:G11652 #:G11400))
 (MFSS (*LIST* #:G5745 #:G11653 #:G11401)) (END-CHUNK 1) (BEGIN-CHUNK 0)
 (STR "gH1 ɛ̄/, #") (BEGIN-PHRASE 0) (END-PHRASE NIL) (CAT NOUN) (PART GA)
 (ROLES ACTOR)))
 (ACTEE
 ((BL +) (INANI-SUBJ +) (COUNT OBJECT) (S #:SUSHI) (RESTRICT +)
 (M (*LIST* SUSI WO)) (STRKEYS (*LIST* "3fW̄]M" "&"))
 (PARTSTRKEYS (*LIST* "M^?W] " "-ɛ")) (ALLKEYS (*LIST* #:G8724 #:G11414))
 (MFSS (*LIST* #:G8725 #:G11415)) (END-CHUNK 2) (BEGIN-CHUNK 2)
 (STR f12FWJɛ") (BEGIN-PHRASE 2) (END-PHRASE NIL) (CAT NOUN) (PART WO)
 (ROLES ACTEE)))
 (PERIOD +)))
```

Figure 13: The raw feature-structure resulting from parsing the sentence above.

6 The Parser

SynfaHir analysis is performed via a neutralized JR parser with pseudo-unification. (The current version of the parser is written in Lisp.) The parser implements a lexical-functional grammar (LFG). However, in designing the syntactic analyzer, we have had to confront several hard problems in NLP—illformed input, ambiguity, and robustness of coverage—which have led us to augment and modify a basic LFG approach.

7 The Grammar

The grammar is a set of about 150 rules written in a pseudo-unification formalism, with some calls to Lisp functions. The grammar covers basic mono-clausal sentences and some sentences involving embedded complement clauses, gerund clauses, and sentential modifiers of nouns. A range of noun phrase structures are also allowed including noun phrases with adjectives, modifiers marked with *no*, and conjuncts marked with *to* and *j/n*. Coverage is sufficient for a first-year Japanese course, with some additional coverage of more complex structures. Figure 1-1 shows a single grammar rule.

```
;;; SURU-VERB
;;; The following two rules are designed to give the same f-structure to verb
;;; phrases with SURU-VERB both with or without "wo".
;;; Nouns that takes the light verb ("suru") is marked by "suru" feature in
;;; the lexicon with the meaning of the verbal,
;;; (eg. "kekkon" s "marriage" suru "marry")
;;; (1) NP SURU
(<VP> <=> (<NP> <VP>) ; "benkyou suru"
  (
    ((x2 suru) =c +) ; verb is "suru"
    ((x1 suru) = *DEFINED*) ; noun takes "suru"
    ((x1 part) = UNDEFINED*)
    (x0=x2)
    ((x0 previousword) <= x1)
    ((x0 previousword separator) <= * I I)
    (x0 <= (combined-word-from-word-and-previousword (x0)))
    ((x0 subcat) <= (x1 subcat)) ;overwrite with the noun's
    ;subcategorization frame
    ((x0 s) <= (x1 suru)) ;overwrite the sense
    ((x0 begin-phrase) <= (x1 begin-phrase))
  ))
```

Figure 1-1: A rule from the parser's grammar.

8 The Mapper

The mapper assigns grammatical functions and/or semantic roles to phrases by determining the best match of phrases with the predicate's subcategorization frame. This is complicated by the fact that student Japanese contains many errors in the use of case marking particles. There is also some ambiguity caused by the use of the topic marker *ira* or the conjunctive particle* *wo* in place of other particles.

The most recent version of the mapper is called from the grammar rules that attach noun phrases to sentences. Each phrase can match zero or more slots in a predicate's subcategorization frame on the basis of its particle and its semantic class. During analysis of student sentences, because of particle errors, there are often phrases that do not match any subcategorization slots. In this case, the mapper searches for the meaning of the phrase in the authored parse, the function/role of the phrase with the same meaning in the authored parse* is hypothesized as the intended function/role in the student sentence, and an error in particle* use* is flagged. Phrases that still do not have a function/role after this procedure are given the function EXTRA. Phrases without particles are given the function NOMINAL. At this point, each phrase is tagged with zero or more slots it can fill. There may be five phrases that fill the same slot, and any given phrase may be able to fill more than one slot. The mapper then identifies one or more ways of filling the predicate's subcategorization frame* possibly by delegating more phrases to the extra function.

The* mapper is written in LISP and TransformationKit, a pseudo-unification language' for specifying transformations of feature structures.

9 The Disambiguator

Because ALK'E-rhan currently has only minimal semantic analysis, there can be many unresolved ambiguities in the out-put of syntactic analysis. The disambiguator resolves syntactic ambiguities interactively with the user. First, it detects certain types of ambiguities in the output of the segmentation program, parser, and mapper. Then it formulates questions which are presented to the user one at a time. Sometimes more than one question is formulated to resolve the same ambiguity, so that the user can answer the one that is most clear. Currently, the disambiguator can detect ambiguities in segmentation, homographs, grammatical function/semantic role assignment, nesting of noun modifiers marked with *no* and conjuncts marked with *lo* and *ya*. and attachment ambiguities involving sentential modifiers of nouns and gerund clauses.

The disambiguator is written in LISP, while the questions are presented and user input is obtained with cT.

Examples of some disambiguation dialogs are shown in Figure 15.

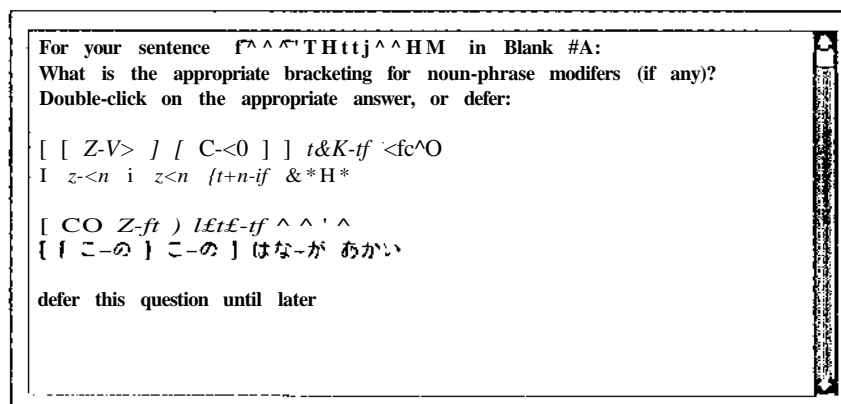



Figure 15: Some disambiguation dialogs.

10 The Matcher

Student responses that are well formed may nevertheless be incorrect in a particular context. For example, they may convey the wrong level of formality or the wrong relationship between speakers, or they may simply express the wrong proposition. For this reason, even when the syntactic analyzer assigns a well-formed structure to a student response, it must be checked for appropriateness in context. This is the job of the matcher.

The matcher compares the analysis of a student's response with the analysis of the teacher's response, which was stored during authoring. Errors are reported if required features or phrases are missing, if extra features or phrases are used, or if the values of any features differ.

Some degree of flexibility is inherent in the matcher because it matches parsed analyses of sentences instead of comparing the sentences themselves. In many cases, different sentences with the same meaning have identical analyses. Notable examples of this in Japanese are sentences that differ in word order, but have the same meaning. If the parser assigns the same analysis to two sentences, the matcher will treat them as identical.¹

The current matcher is written in C as a recursive string-comparison based process.

¹It should be noted, however, that not all synonymous sentences have identical analyses. Our current parser involves increasing the number of types of synonymy that are recognized by the NLP programs.

11 Overview of Error Detection and Feedback

Errors in student responses are detected at three stages. First, during parsing, special grammar rules are designed to parse errorful structures and raise error flags. Second, during mapping, the mapper may detect phrases whose grammatical function or semantic role cannot be determined. Third, during matching, contextually inappropriate features of the student's response are detected by comparing (the analysis of the student's response to Mirrored analysis of the authored response).

causer 私	OPT
HUMAN	REQ
ANIMATE	REQ
particle は	REQ
meaning I	
actor=causee いくお	REQ
PROPER	REQ
HUMAN	REQ
ANIMATE	REQ
particle に	REQ
meaning IKUO--BOY'S FIRST NAME	
thing_acted_on 野菜	REQ
particle を	REQ
meaning VEGETABLES	
predicate 食べさせた	REQ
CAUSATIVE	REQ
meaning EAT	
PAST	REQ
PRESENT	ILL
dictionary-form 食べる	
INFORMAL	OPT
FORMAL	OPT
EXTENDED-PREDICATE	OPT
sentence-particle	OPT

Figure 16: Feature-slot for the author's sentence

```

actor=causee いくお
  HUMAN
  PROPER
  ANIMATE
  particle を
  meaning IKUO--BOY'S FIRST NAME
thing_acted_on 野菜
  particle を
  meaning VEGETABLES
predicate 食べさせました
  meaning EAT
  PAST
  dictionary-form 食べる
  CAUSATIVE
  ERRORPART      PART-ERROR ACTOR=CAUSEE WO いくお
  FORMAL

```

Figure 17: Feature-slot for the student's sentence

In blank # A: You seem to have used the wrong particle on いくお.
 [The particle of the actor=causee (いくお) should be に.]
 You made the predicate (食べさせました) FORMAL, which is optional.

Figure 18: The feedback produced after matching the above two feature-slots
