

JETR: A ROBUST MACHINE TRANSLATION SYSTEM

Rika Yoshii

Department of Information and Computer Science
University of California, Irvine,
Irvine, California, 92717[†]

ABSTRACT

This paper presents an expectation-based Japanese-to-English translation system called JETR which relies on the forward expectation-refinement process to handle ungrammatical sentences in an elegant and efficient manner without relying on the presence of particles and verbs in the source text. JETR uses a chain of result states to perform context analysis for resolving pronoun and object references and filling ellipses. Unlike other knowledge-based systems, JETR attempts to achieve semantic, pragmatic, structural and lexical invariance.

INTRODUCTION

Recently there has been a revitalized interest in machine translation as both a practical engineering problem and a tool to test various Artificial Intelligence (AI) theories. As a result of increased international communication, there exists today a massive Japanese effort in machine translation. However, systems ready for commercialization are still concentrating on syntactic information and are unable to translate syntactically obscure but meaningful sentences. Moreover, many of these systems do not perform context analysis and thus cannot fill ellipses or resolve pronoun references. Knowledge-based systems, on the other hand, tend to discard the syntax of the source text and thus are unable to preserve the syntactic style of the source text. Moreover, these systems concentrate on understanding and thus do not preserve the semantic content of the source text.

An expectation-based approach to Japanese-to-English machine translation is presented. The approach is demonstrated by the JETR system which is designed to translate recipes and instruction booklets. Unlike other Japanese-to-English translation systems, which rely on the presence of particles and main verbs in the source text (AAT 1984, Ibuki 1983, Nitta 1982,

Saino 1983, Shimazu 1983), JETR is designed to translate ungrammatical and abbreviated sentences using semantic and contextual information. Unlike other knowledge-based translation systems (Cullingford 1976, Ishizaki 1983, Schank 1982, Yang 1981), JETR does not view machine translation as a paraphrasing problem. JETR attempts to achieve *semantic, pragmatic, structural and lexical invariance* which (Carbonell 1981) gives as multiple dimensions of quality in the translation process.

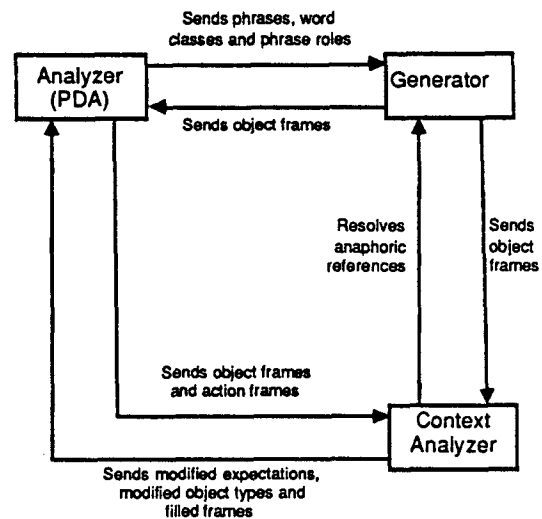


Figure 1. JETR Components

JETR is comprised of three interleaved components: the particle-driven analyzer, the generator, and the context analyzer as shown in Figure 1. The three components interact with one another to preserve information contained in grammatical as well as ungrammatical texts. The overview of each component is presented below. This paper focuses on the particle-driven analyzer.

CHARACTERISTICS OF THE JAPANESE LANGUAGE

The difficulty of translation depends on the similarity between the languages involved. Japanese and English are vastly different languages. Translation from Japanese to English involves restructuring of sentences, disambiguation of words, and additions and

[†]The author is now located at:
Rockwell International Corp.
Autonetics Strategic Systems Division
Mail Code: GA42
3370 Miraloma Avenue, P.O. Box 4192
Anaheim, California 92803-4192

deletions of certain lexical items. The following characteristics of the Japanese language have influenced the design of the JETR system:

1. Japanese is a left-branching, post-positional, subject-object-verb language.
2. Particles and *not* word order are important in determining the roles of the noun phrases in a Japanese sentence.
3. Information is usually more explicitly stated in English than in Japanese. There are no articles (i.e. "a", "an", and "the"). There are no singular and plural forms of nouns. Grammatical sentences can have their subjects and objects missing (i.e. ellipses).

PDA: PARTICLE-DRIVEN ANALYZER

Observe the following sentences:

Verb-deletion:

Neji (screw) o (object marker) migi (right) e (direction marker) 3 kurikku (clicks).

Particle-deletion:

Shio (salt) keiniku (chicken) ni (destination marker) furu (sprinkle).

The first sentence lacks the main verb, while the second sentence lacks the particle after the noun "shio." The role of "shio" must be determined without relying on the particle and the word order.

In addition to the problems of unknown words and unclear or ambiguous interpretation, missing particles and verbs are often found in recipes, instruction booklets and other informal texts posing special problems for machine translation systems. The Particle-Driven Analyzer (PDA) is a robust intrasentence analyzer designed to handle ungrammatical sentences in an elegant and efficient manner.

While analyzers of the English language rely heavily on verb-oriented processing, the existence of particles in the Japanese language and the subject-object-verb word order have led to the PDA's reliance on forward expectations from words other than verbs. The PDA is unique in that it does not rely on the presence of particles and verbs in the source text. To take care of missing particles and verbs, not only verbs but all nouns and adverbs are made to point to *action frames* which are structures used to describe actions. For both grammatical and ungrammatical sentences, the PDA continuously *combines* and *refines* forward expectations from various phrases to *determine their roles* and to *predict actions*. These expectations are semantic in nature and disregard the word order of the sentence. Each expectation is an action-role pair of the form (<action> <role>). Actions are names of

action frames while roles correspond to the slot names of action frames. Since the main verb is almost always found at the end of the sentence, combined forward expectations are strong enough to point to the roles of the nouns and the meaning of the verb. For example, consider "neji (screw) migi (right) e 3 kurikku (clicks)." By the time, "3 clicks" is read, there are strong expectations for the act of turning, and the screw expects to be the object of the act.

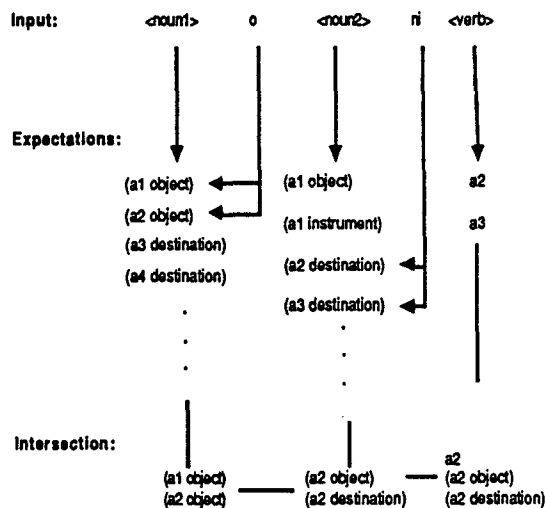


Figure 2. Expectation Refinement in the PDA

Figure 2 describes the forward expectation-refinement process. In order to keep the expectation list to a manageable size, only ten of the most likely roles and actions are attached to each word.

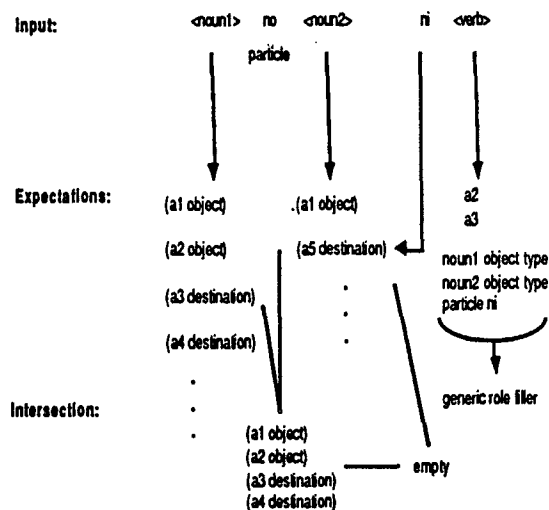


Figure 3. Expectation Mismatch in the PDA

The PDA is similar to IPP (Lebowitz 1983) in that words other than verbs are made to point to structures which describe actions. However, unlike IPP, a generic role-filling process will be invoked *only* if an

unexpected verb is encountered or the forward expectations do not match. Figure 3 shows such a case. The verb will not invoke any role-filling or role-determining process if the semantic expectations from the other phrases match the verb. Therefore, the PDA discourages inefficient verb-initiated backward searches for role-fillers even when particles are missing.

Unlike LUTE (Shimazu 1983), the PDA's generic role-filling process does not rely on the presence of particles. To each slot of each action frame, acceptable filler types are attached. When particles are missing, the role-filling rule matches the object types of role fillers against the information attached to action frames. The object types in each domain are organized in a hierarchy, and frame slots are allowed to point to any level in the hierarchy.

Verbs with multiple meanings are disambiguated by starting out with a set of action frames (e.g. a2 and a3) and discarding a frame if a given phrase cannot fill any slot of the frame.

The PDA's processes can be summarized as follows:

1. Grab a phrase bottom-up using syntactic and semantic word classes. Build an object frame if applicable.
2. Recall all expectations (action-role pairs) attached to the phrase.
3. If a particle follows, use the particle to refine the expectations attached to the phrase.
4. Take the intersection of the old and new expectations.
5. If the intersection is empty, set a flag.
6. If this is a verb phrase and the flag is up, invoke the generic role-filling process.
7. Else if this is the end of a simple sentence, build an action frame using forward expectations.
8. Otherwise go back to Step 1.

To achieve extensibility and flexibility, ideas such as the detachment of control structure from the word level, and the combination of top-down and bottom-up processing have been incorporated.

SIMULTANEOUS GENERATOR

Certain syntactic features of the source text can serve as functionally relevant features of the situation being described in the source text. Preservation of these features often helps the meaning and the nuance to be reproduced. However, knowledge-based systems

discard the syntax of the original text. In other words, the information about the syntactic style of the source text, such as the phrase order and the syntactic classes of the original words, is not found in the internal representation. Furthermore, inferred role fillers, causal connections, and events are generated disregarding the brevity of the original text. For example, the generator built by the Electrotechnical Laboratory of Japan (Ishizaki 1983), which produces Japanese texts from the conceptual representation based on MOPs (Schank 1982), generates a pronoun whenever the same noun is seen the second time. Disregarding the original sentence order, the system determines the order using causal chains. Moreover, the subject and object are often omitted from the target sentence to prevent wordiness.

Unlike other knowledge-based systems, JETR can preserve the syntax of the original text, and it does so without building the source-language tree. The generation algorithm is based on the observation that human translators do not have to wait until the end of the sentence to start translating the sentence. A human translator can start translating phrases as he receives them *one at a time* and can apply partial syntax-transfer rules as soon as he notices a phrase sequence which is ungrammatical in the target language.

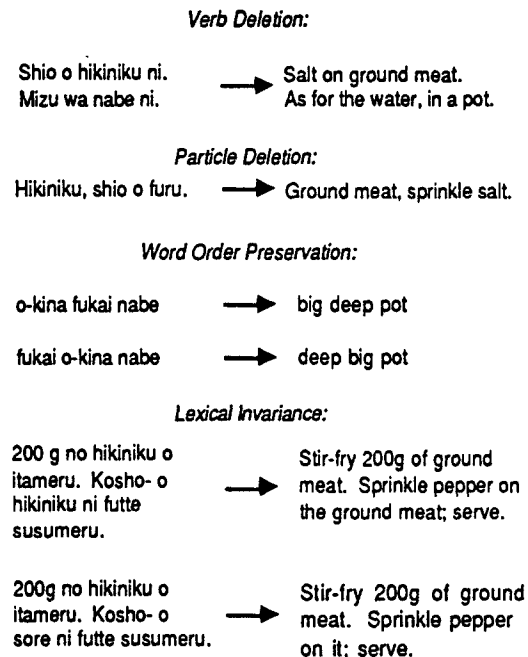


Figure 4. Style Preservation in the Generator

The generator does not go through the complete semantic representation of each sentence built by the other components of the system. As soon as a phrase is processed by the PDA, the generator receives the phrase along with its semantic role and starts generating the phrase if it is unambiguous. Thus the generator can easily distinguish between inferred information and information explicitly present in the

source text. The generator and *not* the PDA calls the context analyzer to obtain missing information that are needed to translate grammatical Japanese sentences into grammatical English sentences. No other inferred information is generated. A preposition is not generated for a phrase which is lacking a particle, and an inferred verb is not generated for a verb-less sentence. Because the generator has access to the actual words in the source phrase, it is able to reproduce frequent occurrences of particular lexical items. And the original word order is preserved as much as possible. Therefore, the generator is able to preserve idiolects, emphases, lengths, ellipses, syntax errors and ambiguities due to missing information. Examples of target sentences for special cases are shown in Figure 4.

To achieve structural invariance, phrases are output as soon as possible without violating the English phrase order. In other words, the generator pretends that incoming phrases are English phrases, and whenever an ungrammatical phrase sequence is detected, the new phrase is saved in one of three queues: SAVED-PREPOSITIONAL, SAVED-REFINER, and SAVED-OBJECT. As long as no violation of the English phrase order is detected or expected, the phrases are generated immediately. Therefore, no source-language tree needs to be constructed, and no structural information needs to be stored in the *semantic* representation of the complete sentence.

To prevent awkwardness, a small knowledge base which relates source language idioms to those of the target language is being used by JETR; however, one problem with the generator is that it concentrates too much on information preservation, and the target sentences are awkward at times. Currently, the system cannot decide when to sacrifice information preservation. Future research should examine the ability of human translators to determine the important aspects of the source text.

INSTRA: THE CONTEXT ANALYZER

The context analyzer component of JETR is called INSTRA (INSTRUction Analyzer). The goal of INSTRA is to aid the other components in the following ways:

1. Keep track of the changes in object types and forward expectations as objects are modified by various modifiers and actions.
2. Resolve pronoun references so that correct English pronouns can be generated and expectations and object types can be associated with pronouns.
3. Resolve object references so that correct expectations and object types can be associated with objects and consequently the article and the number of each noun can be determined.

4. Choose among the multiple interpretations of a sentence produced by the PDA.
5. Fill ellipses when necessary so that well-formed English sentences can be generated.

In knowledge-based systems, the context analyzer is designed with the goal of natural-language understanding in mind; therefore, object and pronoun references are resolved, and ellipses are filled as a by-product of understanding the input text. However, some human translators claim that they do not always understand the texts they translate (Slocum 1985). Moreover, knowledge-based translation systems are less practical than systems based on direct and transfer methods. Wilks (1973) states that "...it may be possible to establish a level of understanding somewhat short of that required for question-answering and other intelligent behaviors." Although identifying the level of understanding required in general by a machine translation system is difficult, the level clearly depends on the languages, the text type and the tasks involved in translation. INSTRA was designed with the goal of identifying the level of understanding required in translating instruction booklets from Japanese to English.

A unique characteristic of instruction booklets is that every action produces a clearly defined resulting state which is a transformed object or a collection of transformed objects that are likely to be referenced by later actions. For example, when salt is dissolved into water, the salty water is the result. When a screw is turned, the screw is the result. When an object is placed into liquid, the object, the liquid, the container that contains the liquid, and everything else in the container are the results. INSTRA keeps a *chain of the resulting states* of the actions. INSTRA's five tasks all deal with searches or modifications of the results in the chain.

```

- ingredients -
OBJ RICEVAMT 3 CUPSALIAS INGO
OBJ WINGVADJ CHICKENVAMT 100 TO 120 GRAMSALIAS ING1
OBJ EGGVAMT 4ALIAS ING2
OBJ BAMBOO:SHOOTVADJ BOILEDVAMT 40 GRAMSALIAS ING3
OBJ ONIONVADJ SMALLVAMT 1ALIAS ING4
OBJ SHIITAKE:MUSHROOMADJ FRESHVAMT 2ALIAS INGS
OBJ LAVERVAMT AN APPROPRIATE AMOUNTALIAS ING6
OBJ MITSUBAVAMT A SMALL AMOUNTALIAS ING7
- the rice is boiled -
STEP1 OBJ RICEVALIAS INGOART TREFPLURAL T
- the chicken, onion, bamboo shoots, mushrooms and mitsuba are cut
STEP2 OBJ CHICKENVALIAS ING1VART TREFPLURAL T
STEP2 OBJ ONIONVALIAS ING4ART T
STEP2 OBJ BAMBOO:SHOOT VALIAS ING3ART TREFPLURAL T
STEP2 OBJ SHIITAKE:MUSHROOMADJ FRESHVALIAS INGSVART T
REFPLURAL T
STEP2 OBJ MITSUBAVALIAS ING7ART T

```

Figure 5. Chain of States Used by INSTRA

To keep track of the state of each object, the object type and expectations of the object are changed whenever certain modifiers are found. Similarly, at the end of each sentence, 1) the object frames representing the result objects are extracted from the frame, 2) each result object is given a unique name, and 3) the type and expectations are changed if necessary and are attached to the unique name. To identify the result of each action, information about what results from the action is attached to each frame. The result objects are added to the end of the chain which may already contain the ingredients or object components. An example of a chain of the resulting states is shown in Figure 5.

In instructions, a pronoun always refers to the result of the previous action. Therefore, for each pronoun reference, the unique name of the object at the end of the chain is returned along with the information about the number (plural or singular) of the object.

For an object reference, INSTRA receives an object frame, the chain is searched backwards for a match, and its unique name and information about its number are returned. INSTRA uses a set of rules that takes into account the characteristics of modifiers in instructions to determine whether two objects match. Object reference is important also in disambiguating item parts. When JETR encounters an item part that needs to be disambiguated, it goes through the chain of results to find the item which has the part and retrieves an appropriate translation equivalent. The system uses additional specialized rules for step number references and divided objects.

Ellipses are filled by searching through the chain backwards for objects whose types are accepted by the corresponding frame slots. To preserve semantic, pragmatic and structural information, ellipses are filled *only* when 1) missing information is needed to generate grammatical target sentences, 2) INSTRA must choose among the multiple interpretations of a sentence produced by the PDA, or 3) the result of an action is needed.

The domain-specific knowledge is stated solely in terms of action frames and object types. INSTRA accomplishes the five tasks 1) without pre-editing and post-editing, 2) without relying on the user except in special cases involving unknown words, and 3) without fully understanding the text. INSTRA assumes that the user is monolingual. Because the method refrains from using inferences in unnecessary cases, the semantic and pragmatic information contained in the source text can be preserved.

CONCLUSIONS

This paper has presented a robust expectation-based approach to machine translation which does not view machine translation as a testbed for AI. The paper has shown the need to consider problems unique to machine translation such as preservation of syntactic

and semantic information contained in grammatical as well as ungrammatical sentences.

The integration of the forward expectation-refinement process, the interleaved generation technique and the state-change-based processing has led to the construction of an extensible, flexible and efficient system. Although JETR is designed to translate instruction booklets, the general algorithm used by the analyzer and the generator are applicable to other kinds of text. JETR is written in UCI LISP on a DEC system 20/20. The control structure consists of roughly 5500 lines of code. On the average it takes only 1 CPU second to process a simple sentence. JETR has successfully translated published recipes taken from (Ishikawa 1975, Murakami 1978) and an instruction booklet accompanying the Hybrid-H239 watch (Hybrid) in addition to hundreds of test texts. Currently the dictionary and the knowledge base are being extended to translate more texts.

Sample translations produced by JETR are found in the appendix at the end of the paper.

REFERENCES

- AAT. 1984. Fujitsu has 2-way Translation System. AAT Report 66. Advanced American Technology, Los Angeles, California.
- Carbonell, J. G.; Cullingford, R. E. and Gershman, A. G. 1981. Steps Toward Knowledge-Based Machine Translation. *IEEE Transaction on Pattern Analysis and Machine Intelligence* PAMI, 3(4).
- Cullingford, R. E. 1976. The Application of Script-Based Knowledge in an Integrated Story Understanding System. *Proceedings of COLING-1976*.
- Granger, R.; Meyers, A.; Yoshii, R. and Taylor, G. 1983. An Extensible Natural Language Understanding System. *Proceedings of the Artificial Intelligence Conference, Oakland University, Rochester, Michigan*.
- Hybrid. Hybrid-cal. H239 Watch Instruction Booklet. Seiko, Tokyo, Japan.
- Ibuki, J; et. al. 1983. Japanese-to-English Title Translation System, TITRAN - Its Outline and the Handling of Special Expressions in Titles. *Journal of Information Processing*, 6(4): 231-238.
- Ishikawa, K. 1975. *Wakamuki Hyoban Okazu 100 Sen*. Shufu no Tomo, Tokyo, Japan.
- Ishizaki, S. 1983. Generation of Japanese Sentences from Conceptual Representation. *Proceedings of IJCAI-1983*.
- Lebowitz, M. 1983. Memory-Based Parsing. *Artificial Intelligence*, 21: 363-404.
- Murakami, A. 1978. *Futari no Ryori to Kondate*. Shufu no Tomo, Tokyo, Japan.
- Nitta, H. 1982. A Heuristic Approach to English-into-Japanese Machine Translation. *Proceedings of COLING-1982*.

- Saino, T. 1983. Jitsuyoka e Ririku Suru Shizengengo Shori-Gijutsu. *Nikkei Computer*, 39: 55-75.
- Schank, R. C. and Lytinen, S. 1982. Representation and Translation. Research Report 234. Yale University, New Haven, Connecticut.
- Shimazu, A; Naito, A. and Nomura, H. 1983. Japanese Language Semantic Analyzer Based on an Extended Case Frame Model. Proceedings of IJCAI-1983.
- Slocum, J. 1985. A Survey of Machine Translation: Its History, Current Status and Future Prospects. *Computational Linguistics*, 11(1): 1-17.
- Wilks, Y. 1973. An Artificial Intelligence Approach to Machine Translation. In: Schank, R. C. and Colby, K., Eds., *Computer Models of Thought and Language*. W. H. Freeman, San Francisco, California: 114-151.
- Yang, C. J. 1981. High Level Memory Structures and Text Coherence in Translation. Proceedings of IJCAI-1981.
- Yoshii, R. 1986. JETR: A Robust Machine Translation System. Doctoral dissertation, University of California, Irvine, California.

APPENDIX - EXAMPLES

NOTE: Comments are surrounded by angle brackets.

EXAMPLE 1

SOURCE TEXT: (Hybrid)

Anarogu bu no jikoku:awase.

60 pun shu-sei.

Ryu-zu o hikidashite miagi e subayaku 2 kurikku mawasu to cho-shin ga 1 kaiten shite 60 pun susumu. Mata gyaku ni, hidari e subayaku 2 kurikku mawasu to cho-shin ga 1 kaiten shite 60 pun modoru. Ryu-zu o 1 kurikku mawasu tabigoto ni pitt to iu kakuninon ga deru.

TARGET TEXT:

The time setting of the analogue part.

The 60 minute adjustment.

Pull out the crown; when you quickly turn it clockwise 2 clicks, the minute hand turns one cycle and advances 60 minutes. Also conversely, when you quickly turn it counterclockwise 2 clicks, the minute hand turns one cycle and goes back 60 minutes. Everytime you turn the crown 1 click, the confirmation alarm "peep" goes off.

EXAMPLE 2

SOURCE TEXT: (Murakami 1978)

Tori no karaage.

4 ninmae.

<<ingredients need not be separated by punctuation>>
honetsuki butsugiri no keiniku 500 guramu
jagaimo 2 ko
kyabetsu 2 mai
tamanegi 1/2 ko
remon 1/2 ko
paseri.

(1).
Keiniku ni sho-yu o-saji 2 o karamete 1 jikan oku.

(2).
Jagaimo wa yatsuwari ni shite kara kawa o muki mizu ni 10 pun hodo sarasu. <<wa is an ambiguous particle>>

(3).
Tamanegi wa usugiri ni shite mizu ni sarashi kyabetsu wa katai tokoro o sogitotte tate ni 3 to-bun shite kara hosoku kizami mizu ni sarasu.

(4).
Chu-ka:nabe ni abura o 6 bunme hodo irete chu-bi ni kakeru.

(5).
Betsu nabe ni yu o wakashi jagaimo no mizuke o kitte 2 fun hodo yude zaru ni agete mizuke o kiru.

(6).
(1) no keiniku no shiruke o kitte komugiko o usuku mabusu.

(7).
Jagaimo ga atsui uchini ko-on no abura ni ire ukiagatte kita ra chu-bi ni shi kitsuneiro ni irozuite kita ra tsuyobi ni shite kararito sasete ageami de tebayaku sukuiage agedai ni totte abura o kiru.

(8).
Keiniku o abura ni ire ukiagatte kita ra yowame no chu-bi ni shite 2 fun hodo kakete naka made hi o to-shi tsuyobi ni shite kitsuneiro ni ageru. <<hi o to-shi is idiomatic>>

(9).
(3) no tamanegi, kyabetsu no mizuke o kiru. Kyabetsu o utsuwa ni shiite keiniku o mori jagaimo to tamanegi o soe remon to paseri o ashirau.

TARGET TEXT:

Fried chicken.

4 servings.

500 grams of chopped chicken
2 potatoes
2 leaves of cabbage
1/2 onion
1/2 lemon
parsely

(1).
All over the chicken place 2 tablespoons of soy sauce; let alone 1 hour.

(2).

As for the potatoes, after you cut them into eight pieces, remove the skin; place about 10 minutes in water.

(3).

As for the onion, cut into thin slices; place in water. As for the cabbage, remove the hard part; after you cut them vertically into 3 equal pieces, cut into fine pieces; place in water.

(4).

In a wok, place oil about 6/10 full; put over medium heat.

(5).

In a different pot, boil hot water; remove the moisture of the potatoes; boil about 2 minutes; remove to a bamboo basket; remove the moisture.

(6).

Remove the moisture of the chicken of (1); sprinkle flour lightly.

(7).

While the potatoes are hot, place in the hot oil; when they float up, switch to medium heat; when they turn golden brown, switch to strong heat; make them crispy; with a lifter drainer, scoop up quickly; remove to a basket; remove the oil.

(8).

Place the chicken in the oil; when they float up, switch to low medium heat; put over the heat about 2 minutes; completely let the heat work through; switch to strong heat; fry golden brown.

(9).

Remove the moisture of the onion of (3) and the cabbage of (3); spread the cabbage on a dish; serve the chicken; add the potatoes and the onion; add the lemon and the parsely to garnish the dish.