

Generating Natural Sentences by Using Shallow Discourse Information

Shiho Ogino

Tetsuya Nasukawa

IBM Research, Tokyo Research Laboratory
1623-14, Shimotsuruma, Yamato-shi, Kanagawa-ken 242, JAPAN
{shiho, nasukawa}@trl.ibm.co.jp

Abstract. One of the biggest defects of natural language generation systems is that the output sentences are unnatural and contain many redundancies. Machine translation (MT) users, for instance, often get tired of reading the output of MT because of this problem. In this paper, we summarize the results of our analysis of human translation in terms of the use of discourse information to generate target-language sentences, and describe our attempt to generate natural output by referring to a simple discourse model that consists of a syntactic parser's output for each sentence in the discourse.

1 Introduction

Sentence generation is a key component of natural language processing systems in the sense that it is the final stage of the process for producing system outputs, using all the information processed in the entire system, and also in the sense that improvement of the component directly results in improvement of the system output. However, in the domain of MT systems, much effort has been devoted to generating accurate sentences that are often redundant and unnatural for human readers. This paper describes our attempt to generate natural sentences, based on studies of translation by human translators.

To avoid redundancy in texts, humans often use anaphoric expressions, omissions, rewordings, and so on. Much research in natural language processing has focused on analyzing such phenomena [Hobbs, 1978, Lappin & McCord, 1990, Nakaiwa & Ikehara, 1992] in order to arrive at a correct interpretation of the text. Such analysis often requires discourse information [Schank & Riesbeck, 1981, Grosz & Sidner, 1986], since the antecedent of an expression may not appear within the same sentence. On the other hand, the results of discourse processing are rarely reflected explicitly in sentences generated by practical NLP systems. There are two main reasons for this: one is that they may lead to misinterpretation caused by misresolution of pronoun referents [Nasukawa, 1996]; the other is that discourse processing systems often require an enormous amount of background knowledge and are therefore considered high-cost items.

However, within a restricted domain such as that of computer manuals, many context-dependent problems seem solvable without the use of a deep inference mechanism. We tried to develop a practical method that requires only a fairly simple and shallow discourse model and solves most context-dependent problems. To begin with, we analyze human translation in terms of the use of discourse information to generate target language sentences, because human translation offers a good example of how to generate natural sentences. We also summarize how to generate natural outputs by referring to a simple discourse model that consists of a syntactic parser's output for each sentence in a discourse. As data, we used 1,012 English sentences from a computer manual and the same sentences translated into Japanese by a human translator. We focus on how pronouns and determiners are translated, because this is a typical context-dependent problem.

In Section 2, we briefly introduce our method of shallow discourse processing and describe what is required in natural sentence generation and what can be achieved by using the results of shallow discourse processing. Section 3 describes the results of our analysis of human translation in terms of the use of discourse information. The results of our attempt to apply our method to machine translation output are discussed in Section 4.

2 Shallow discourse processing applied to translation

To refer to context information that consists of data on one or more sentences in a text, it is essential to construct some context model, even if only shallow information will be used

in the analysis process. We use only a set of parsed trees (results of syntactic analysis) as a context model; neither a huge amount of background knowledge nor a precise discourse model is necessary for our approach. In machine translation, the parse trees are obtained through syntactic parsing [Nasukawa, 1996], but in other types of processing they can be created by the concept builder or the spoken-language-processing module. For shallow discourse analysis, it is sufficient to obtain a set of parse trees as input, no matter how they are created. The following information must be stored in each parse tree:

- Surface string of each phrase
- Stem of the head of each phrase
- Part of speech of the head of each phrase
- Modification relation of each phrase
- Case roles of the noun phrases with respect to their governor verb phrase

In addition, if any on-line knowledge resources are available, information extracted from the resources is also stored in the context model. For example, in our machine translation system, information on synonyms extracted from an on-line thesaurus [Collins, 1984] is added to the context model. However, except for this existing information, no other information specially hand-coded for context processing is added to the context model.

In the second step, each sentence in the context model is analyzed individually, and its ambiguities and context-dependent problems are resolved by referring to information on other sentences in the context model. Pronoun resolution is one of these context-dependent problems that must be resolved.

Our context model is used to select candidate noun phrases for a pronoun referent. Information on word frequency and collocation patterns extracted from the context model improves the accuracy with which the correct referent is selected from the candidate noun phrases. We have now achieved a success rate of 93.8% in pronoun resolution [Nasukawa, 1994].

Noun phrases with determiners are similar to pronouns in the sense that they often have antecedents in the text. Resolution of the antecedents of noun phrases with determiners can be considered useful for generating natural sentences. We used the procedure for finding pronoun referents to manually mark to the candidate antecedents of each noun phrase with a determiner. Finally, the nearest candidate in the previous context is considered to be the antecedent of that noun phrase.

When some sentence conjunctions, such as "however," "but," and "and," are placed at the head of a sentence, they explicitly expose the sentence's relation to a set of other sentences. We can determine some sentence-to-sentence relations by considering these conjunctions as relation markers. In the domain of computer manuals, less complex structures are normally used to mark sentence-to-sentence relationships, but the above markers can also be considered in this domain, since they expose a sentence's relation to the immediately preceding sentence. They can be also treated as signs of sentence concatenation, as we will explain in the next section.

3 Analysis of human translation

3.1 Translation patterns for Japanese manuals

There are many standard rules for translating manuals. For instance, "you" should be translated as "yuuza" (user) in English-to-Japanese translation. In the data sentences, we can find the following rules:

- "You" should be omitted if possible.
- "You" should be translated as "yuuza" (user) if it cannot be omitted.
- "This manual" should be translated as the compound noun "honsho" (this book).

When an English expression matches one of the patterns above, human translators usually adopt the corresponding rule in order to make the translation style even.

3.2 Pronoun

In Japanese, an English pronoun is often translated as a pronoun, as its referent, or as a sequence of determiner and referent. To replace a pronoun with its referent, it is necessary to use discourse information. We observed what proportion of pronouns are translated by using

High-Speed-Capture-Subsystem Set Status

This is the status flag for this set in the VSAM RPDS file.
High-Speed-Capture-Subsystem Scanner Data Field 5
This is a data field extracted from the image by the scanner.

Figure 1. Definition list

discourse information in human translation in order to estimate how our discourse processing affects the generation of sentences.

It is sometimes said that it is generally preferable to generate zero-pronouns rather than real pronouns in Japanese, because some "indispensable" phrases (subject, direct object, indirect object, and so on) can be omitted in Japanese, especially when their referents appear repeatedly in the context. The second part of this statement may be right in many cases, but our analysis of human translation shows that the first part is not adequate in generating written Japanese.

The 1,012 English sentences that we used as data contained one hundred and nine pronouns. Of these, sixty-four were "you." In the translation, "you" was omitted in forty-nine out of these sixty-four cases. In other words, 76.5% of all occurrences of "you" were translated by using the manual translation rules mentioned in Subsection 3.1. Other pronouns, however, were omitted in only twenty-three out of forty-five cases in translation; that is, in half of all occurrences, pronouns other than "you" were translated without omission. In addition, sixteen of the twenty-three omitted pronouns appeared not in ordinary sentences but in definition lists of technical terms. The sentence set in Figure 1 is an example of a definition list. Omitted pronouns are written in italic letters.

In these definition lists, translating every instance of "this" into Japanese for each definition often leads to redundancy, especially when the list is quite long, because "kore" or "sore," the Japanese pronouns corresponding to "this," might be repeated throughout the list. To avoid such redundancy, there seems to be another rule for manual translations:

Pronouns in definition lists should be omitted if possible in generating Japanese sentences.

Of the twenty-two pronouns that were translated without omission, the antecedents of seven could be resolved by using inter-sentential information, and the pronouns were replaced with their referents in the generated output, while five were translated by using intra-sentential information, and eight were translated by using equivalent Japanese pronouns. That is, inter-sentential information is useful for translating about 30% of pronouns that should not be omitted.

Four out of the last-mentioned eight pronouns had their antecedents in the same sentence, while two had antecedents in the previous context.

Two of the twenty-two translated pronouns were translated by using knowledge that could not be acquired from the source text. Table 1 shows details of the pronoun translation methods.

3.3 Determiner

In English-to-Japanese machine translation, most English determiners are omitted, because there are no equivalent words in Japanese. In human translation, however, some are rendered as the corresponding Japanese determiners (e.g. "kore" and "sore"). We assume that the differences in the way determiners are translated are determined by the existence of antecedents. To confirm

Table 1. Translation of pronouns

	A	B	C	D	
	a	b	c	d	e
"You"	0	2	49	0	0
Other pronouns	8	-	23	5	7
Total	8	2	72	5	7
				15	109

A Translate without discourse processing
B Either A or C
C Translate by using shallow discourse information
D Other
a Translate by using an equivalent pronoun
b Translate according to a rule for manuals
c Omit
d Translate by using the referent
e Translate by using the referent with a determiner
f Other

Processing Overview

The HSCS is activated when the operator initiates a session using the OSCN transaction.

The HSCS returns a BEGIN-BATCHER command to the Imageplus workstation.

The BEGIN-BATCHER command has a return code set to non-zero if batch processing cannot be started at this time.

Figure 2. Examples of "the"

this assumption, and to estimate how discourse processing can be applied in generating and omitting Japanese determiners, we investigated how English determiners are translated into Japanese in terms of whether they have antecedents. Tables 2 and 3 show the results of this investigation.

"The" has different properties from other pronoun-form determiners ("this," "that," "these," "those," and so on); no pronoun-form determiners are omitted in the translation of the data sentences, while "the" is omitted far more frequently. This difference may be caused by the fact that pronoun-form determiners in English almost always have antecedents, as the corresponding determiners in Japanese do, while "the" often has no antecedents. Since the properties of English pronoun-form determiners and Japanese determiners are similar, human translators of manuals tend to translate English pronoun-form determiners directly into equivalent Japanese words, in order to avoid misunderstanding. Table 2 shows that discourse information can be useful in only a few cases in translating pronoun-form determiners that appear in the data sentences.

We classified all instances of the English determiner "the" in the data text into the following four categories, and investigated how their translations differ according to their categories.

A "The" appears in an idiomatic phrase (e.g. the following, the first or the biggest): 94

B "The" appears with a proper noun (e.g. the HSCS or the center of Tokyo): 247

C "The" appears in a noun phrase that has no antecedent in the text: 177

D "The" appears in a noun phrase that has an antecedent in the text: 74

The sentences in Figure 2 are examples of the groups C and D. The underlined determiner is an example of C; the determiner written in italics is an example of D.

The noun phrase "the operator" is classified as belonging to group C because it appears for the first time in the text and has no antecedent, while the phrase "the BEGIN-BATCHER command" has an antecedent phrase in the previous sentence and is classified as belonging to group D.

The determiners in group B should be omitted when Japanese sentences are generated. Phrases that include any of the determiners in group A should be translated by using some fixed Japanese word sequences (e.g. "the following X" should be translated by using the sequence "tsugi no" or "ika no"). The determiners in group C are usually omitted in machine translation, because the corresponding determiners in Japanese require antecedent phrases. In some cases, however, determiners classified as belonging to group C are translated by using corresponding Japanese determiners, because although their antecedents do not actually appear in the previous context, they exist in the deep semantic meaning. For instance, "Value" does not appear in the sentence set in Figure 3 in which the underlined phrase "the value" occurs, but the translation of this phrase includes a Japanese determiner, "kono", corresponding to "the."

In this case, the Japanese determiner is generated because the antecedent of "the value" cannot be found in the previous context, though it exists at a deep semantic level [Charniak, 1973]. To solve this problem, we need not only shallow discourse processing but also deep semantic

Table 2. Translation of pronoun-form determiners

A	B	C				
a	b	c	d	e	f	Total
75	12	1	0	3	6	97

A Translate without discourse processing
 B Translate by using shallow discourse information
 C Other
 a Translate by using an equivalent determiner
 b Translate according to a rule
 c Translate by using the antecedent
 d Omit the determiner
 e Omit the noun phrase including the determiner
 f Other

High-Speed-Capture Subsystem LU 6.2 Type

Identifies the type of Image Plus workstation.

E: The value must be B (High-Speed Scanner).

J: Kono (this) atai-wa (value) B-de (B) nakerebanaranai (must).

Figure 3. Japanese determiner corresponding to "the"

analysis, including more complex discourse processing.

Nine out of seventy-four determiners classified as belonging to group D were translated by using corresponding determiners in Japanese. In eight out of nine cases, the distance between these determiners and their antecedents were less than five sentences. This leads us to the assumption that the range within which antecedent phrases of these determiners are likely to occur is narrower than for antecedents of pronouns, because the corresponding Japanese determiners also function as pronoun-form determiners that require their antecedents to lie within narrower range. This result of the investigation also indicates that if the antecedent of a noun phrase containing "the" can be found in the same sentence or in the four preceding English sentences, the probability of the corresponding determiners being generated will increase when the sentences are translated into Japanese.

Another reason for this narrow range is that the structure of sentences in computer manuals is not complex; that is to say, the range in newspaper editorials or novels may be wider. The width of the range may also vary according to the language; the difference in range between the source (or original) concept and the target language to be generated should be considered in each instance of language generation.

3.4 Sentence concatenation and ellipsis

In human translation, it often happens that two or more sentences in the source language are translated as a single sentence in the target language. We investigated the major factor reasons for sentence concatenation in human translation, and cases in which sentence concatenation can be performed by using the results of discourse processing.

In the data sentences, There are twelve cases in which two English sentences are translated into one Japanese sentence. Most of them require deep knowledge or precise context models for generation, but four out of twelve can be solved by using the results of shallow discourse processing, as follows:

If the current sentence has one of the sentential conjunctions "however," "but," "and," or "then" indicating its relation to the previous sentence, concatenate the current-sentence with the previous sentence by using the corresponding conjugative particle (setsuzoku-joshi in Japanese) in Japanese generation.

This transformation occurs because frequent use of sentential conjugation is avoided in Japanese sentence style. In generating sentences in other language, the solution will differ according to the preferred style in that language.

When two noun phrases in a sentence that are resolved as coreferential play the same role in relation to the different verb phrases that are sisters syntactically, one of the noun phrases can be omitted in generating a Japanese sentence. In other words, one of the noun phrases can be changed into a zero-pronoun. In English, it always appears as a real pronoun (Figure 4) if it is not the subject of each verb phrase.

This transformation occurs very often when English pronouns are translated into Japanese. Other noun phrases, on the other hand, are not ellipped so often, even if the noun phrases are resolved as coreferential in manual translation, in order to avoid misunderstanding. We can apply the above mentioned procedure for omitting a noun phrase not only to concatenating sentences, but also to translating English pronouns and noun phrases with determiners into Japanese, by using the results of shallow discourse processing.

Table 3. Translation of "the"

a	b	c	d	Total
424	9	94	65	592

a Omit

b Translate by using a corresponding Japanese determiner

c Idiomatic phrase

d Other

Mary buys an apple. Mary eats the apple.
 E: Mary buys an apple and eats it.
 J: Mary-wa (Mary) ringo-wo (apple:obj) katte (buys)₋ (zero-pronoun) taberu (eat).

Figure 4. Ellipsis of noun phrases

4 Discussion

4.1 What information can be referred to in order to generate natural sentences?

As we mentioned in Section 2, the following types of elementary information can be used in sentence generation, since they are assumed to be acquired as the result of parsing the source sentences:

surface string of each phrase, stem of the head of each phrase, part of speech of the head of each phrase, modification relation of each phrase, case roles of the noun phrases with respect to their governor verb phrase.

In addition, we can use the following information acquired as the result of shallow discourse processing [Nasukawa, 1994]:

- Resolved antecedent of each pronoun
The value is set to null when the discourse processing module cannot find an antecedent.
- Resolved antecedent of each noun phrase with determiners
The value is set to null when the discourse processing module cannot find an antecedent.
- Marked conjunction that explicitly exposes a sentence-to-sentence relation.

In the next subsection, we propose a procedure for generating natural sentences by using this information.

4.2 Generation procedure using shallow discourse information

By using the results of our analysis of human translation, we devised a natural Japanese generation procedure for English-to-Japanese translation that should be executed before the conventional generation procedure, as follows:

1. For each sentence, perform the following step:
If the current sentence has one of the sentential conjunctions "however," "but," "and," or "then" indicating its relation to the previous sentence, concatenate the current sentence with the previous sentence by using the corresponding conjugative particle. If not, go to the next step.
2. For each pronoun, perform the following steps:
 - (a) If the current pronoun appears in a word pattern that should be translated according to a translation standard, adopt the rule to be applied. If not, go to the next step.
 - (b) If the current pronoun appears in the main clause, and a noun phrase that is resolved as coreferential with the current pronoun plays the same role in relation to a verb phrase that is a sister of the verb phrase directly governing the current pronoun, omit the current pronoun. If not, go to the next step.
 - (c) If the current pronoun appears in a subordinate clause, and a noun phrase that is resolved as coreferential with the current pronoun plays the same role in relation to a verb phrase that is a sister of the verb phrase directly governing the current pronoun, replace the current pronoun with its antecedent and omit the noun phrase resolved as coreferential. If not, go to the next step.
 - (d) If the current pronoun appears in a definition sentence in a definition list, perform the following step:
If the current pronoun is resolved as coreferential with the entry of the current definition and functions as subject of the current sentence, omit it. If not, go to the next step.
 - (e) If the current pronoun has not yet been omitted and an antecedent noun phrase exists in the previous context, replace the pronoun with its antecedent.
3. For each noun phrase containing "the," perform the following steps:
 - (a) Perform all the above steps except the last one.
 - (b) If the head of the current noun phrase is a proper noun, skip to the next step. If not, perform the following steps:

Procedures for restoring the HSCS journal files are also provided by the system programmer. *These* are recovery procedures and should only be executed at the request of the system programmer.

Human: Korera-wa (these) kaifuku-tetsuzuki-deari (be recovery procedures) ...

Output: Korera-no (these) tetsuzuki-wa (procedures) kaifuku-tetsuzuki-deari (be recovery procedures) ...

Figure 5. Human translation and discourse processing

- i. If the antecedent of the current noun phrase has the determiner "the," and no antecedent of that antecedent noun phrase exists, skip the following step.
 - ii. If the antecedent of the current noun phrase appears in the current sentence or in the immediately preceding four sentences, and the head of that antecedent is not a pronoun, select the corresponding Japanese determiner as the translation of the current determiner. If not, delete the current determiner.
4. Generate Japanese sentences according to the existing algorithm.

This procedure sometimes produces different results from the human translation result. We discuss the differences in the next subsection.

4.3 Result of applying the procedure

In using the procedure to translate pronouns, it is difficult to decide whether the pronoun should be translated by using only the antecedent or the sequence of the corresponding determiner and the antecedent. According to the procedure mentioned above, pronouns are always translated by using the antecedent noun phrases with the corresponding determiners. The results of this process differ from the results of human translation if the human translator uses the corresponding Japanese pronouns or antecedent noun phrases without determiners. Figure 5 shows an example of this difference. The result of the generation procedure in the figure is natural enough to read. In these cases, our shallow discourse processing causes no misreading and often as a good effect on the sentence generation. The same thing can be said in the case where an English pronoun is translated by using the corresponding Japanese pronoun. Most of the cases in which a human translator translates an English pronoun by using the corresponding Japanese pronoun occur when the antecedent of the source pronoun is not a noun phrase (e.g. when it is a full sentence, a set of sentences, or some figure appearing on the same page). Our procedure also outputs the corresponding Japanese pronoun in such cases. When the antecedent can be resolved and lies within the specified range, it is not unnatural in most cases to translate the source pronoun by using a sequence of the corresponding determiner and the antecedent noun phrase.

It is also difficult to decide whether the determiner "the" should be translated by using the corresponding Japanese determiner or omitted. When our procedure omits the source determiner for which the human translator generates the corresponding determiner, it causes no serious problem; it produces a somewhat stiff sentence, but causes no misreading. However, when our procedure produces an extra determiner, it sometimes causes a misreading. An example is shown in Figure 6. Suppose that the phrase in italics, "CICS journal," is resolved as the antecedent of the underlined phrase "the journal." "CICS journal" and "the journal" can share the same semantic information, but "the journal" is not exactly equivalent to the journal that the HSCS stores messages in; it could indicate a journal in which the HSCS has not yet stored messages. In this sense, "CICS journal" and "the journal" denote different things. However, if the extra determiner is generated in the target sentence as it is in the sample, it may mislead the reader into believing that "CICS journal" and "the journal" denote the same thing. In such cases, our procedure has a bad effect on the generation. When we applied the procedure, the generation of an extra determiner that could cause misreading occurred four times in processing seventy-four determiners whose antecedents existed in the previous context. To avoid this side-effect, deeper analysis of discourse information is required.

If the above mentioned points are taken into account, twenty-eight out of forty-five pronouns (53%) other than "you" can be translated naturally by using the results of shallow discourse processing. This rate includes omitted pronouns and pronouns translated by using the corresponding Japanese words. As for determiners, nine out of seventy-four instances of "the" (12%) can be translated naturally by using the results of shallow discourse processing; on the other

The HSCS stores HSCS messages in a *CICS journal*.

The system programmer can provide the number of the journal used by the HSCS.

Human: Shisutemu-purograma-wa (system programmer) HSCS-de (by the HSCS) shiyosuru (used) jaanaru-no (the journal) bango-wo (the number of) shiteisuru (provide) kotoga-dekiru (can).

Procedure: Shisutemu-purogurama-wa (system programmer) HSCS-de (by the HSCS) shiyosuru (used) sono jaanaru-no (that journal) bango-wo (the number of) shiteisuru (provide) kotoga-dekiru (can).

Figure 6. Extra determiner

hand, four out of seventy-four (5%) suffer from the side-effects of shallow discourse processing. This error rate in applying the procedure to translating determiners is too high in relation to the successful application rate; we must consider a new way of dealing with this side-effect. Out of all instances of "the," five hundred and twenty-three out of five hundred and ninety-two (88%) can be translated naturally.

5 Summary

Human translation is a good example of natural sentence generation. We have attempted to determine what information is required for natural language generation by analyzing human translation, and what kinds of information required for generating natural sentences can be acquired by using shallow discourse processing. We have also proposed a procedure for natural sentence generation, using the results of our analysis. Application of the procedure shows that quite shallow discourse processing with a simple discourse model provides rich information for generating natural sentences. The advantage of this method is that it requires neither a huge amount of background knowledge nor a large syntactically tagged corpus.

We plan to extend the application of discourse processing to other domains and to demonstrate how our simple discourse processing method can be applied to various areas of natural language processing.

References

- [Charniak, 1973] E. Charniak. 1973. Jack and Janet in Search of a Theory of Knowledge. In *Advanced Papers from the Third International Joint Conference on Artificial Intelligence*, Stanford, California, pages 337-343.
- [Grosz & Sidner, 1986] B. J. Grosz and C. L. Sidner. 1986. Attentions, Intentions, and the Structure of discourse. *Computational Linguistics*, 12, pages 175-204.
- [Hobbs, 1978] J. R. Hobbs. 1978. Resolving Pronoun References. *Lingua*, 44, pages 311-338.
- [Lappin & McCord, 1990] S. Lappin and M. McCord. 1990. Anaphora Resolution in Slot Grammar. *Computational Linguistics*, 16, pages 197-212.
- [Nakaiwa & Ikehara, 1992] H. Nakaiwa and S. Ikehara. 1992. Zero Pronoun Resolution in a Japanese to English Machine Translation System using Verbal Semantic Attributes. In *Proceedings of the 4th Conference on Applied Natural Language Processing*, Trento, Italy, pages 201-208.
- [Nasukawa, 1994] T. Nasukawa. 1994. Robust Method of Pronoun Resolution Using Full-Text Information. In *Proceedings of COLING-94: The 15th International Conference on Computational Linguistics*, Kyoto, Japan, pages 1157-1163.
- [Nasukawa, 1996] T. Nasukawa. 1996. Full-text processing: improving a practical NLP system based on surface information within the context. In *Proceedings of COLING-96: The 16th International Conference on Computational Linguistics*, Copenhagen, Denmark, pages 824-829.
- [Schank & Riesbeck, 1981] R. C. Schank and C. K. Riesbeck. 1981. *Inside Computer Understanding: Five Programs plus Miniatures*. Lawrence Erlbaum Associates.
- [Shirai et al., 1993] S. Shirai, S. Ikehara and T. Kawaoka. 1993. Effects of Automatic Rewriting of Source Language within a Japanese to English MT System. In *Proceedings of the Fifth Conference on Theoretical and Methodological Issues in Machine Translation*, Kyoto, Japan, pages 226-238.
- [Collins, 1984] Collins Publishers. 1984. *The New Collins Thesaurus*, William Collins Sons & Co. Ltd., Glasgow.
- [Quirk et al., 1985] R. Quirk, S. Greenbaum, G. Leech, and J. Svartvik. 1984. *A Comprehensive Grammar of the English Language*. Longman Group Ltd., London.