7 Generation

In this chapter we shall discuss the **generation** of target language texts. We should immediately distinguish this use of the term 'generation' from its use in phrases like 'first generation' and 'second generation' which refer to historical stages in the development of MT systems (Chapters 1 and 4). An alternative term for the computation of target texts from intermediate representations is **synthesis**. Some authors have tried to make a distinction between 'generation' and 'synthesis', depending on the degree of specification of the input to this phase: 'synthesis' would involve more interpretation of the input structure than more straightforward 'generation'. In this and following chapters, however, we shall continue to treat the two terms as synonymous, and, generally, prefer the more common term 'generation'.

7.1 Generation in direct systems

We begin by looking at the notion of target language generation in direct systems of the 'first generation'. From an examination of the basic architecture of direct systems (see Figure 4.1, reproduced here as Figure 7.1) and from the description of transfer processes in these systems in the last chapter (section 6.4), it should already be clear that in a sense there is no interesting target-language generation at all, as such. In comparison with the 'indirect' interlingual and transfer-based systems (Figures 4.2 and 4.4) it is hard to say where analysis of the source language stops and generation of the target language starts: there is processing oriented towards the target language at an early stage (lexical substitution during

Figure 7.1 Direct MT system

dictionary look-up) and there is processing influenced by the source language at late stages (local reordering of structures).

The 'local reordering' phase can be seen as a mixture of transfer and generation. Input to the module is target language lexical items in source language word order. It is a 'transfer' process in so far as source text structures are transformed into target text structures (cf. Chapter 6), and it represents a 'generation' process in so far as the output consists of target language lexical items in target language sequences. An example of 'local reordering' was given in section 6.4 involving the reversal of the adjective–noun sequence in translating from English into French. The final phase ('text synthesis') is concerned, therefore, almost exclusively with procedures to ensure correct word endings (e.g. *try* plus *-ed* as *tried*) and morphological agreement (e.g. adjectives and nouns in French). It is the only stage which can be rightly called one of 'generation' in that it involves only target language information and operates independently of the source text. Nevertheless, it is rudimentary, consisting only of morphological processing.

In direct systems, therefore, generation is based as much as possible on source language structures: nothing is changed more than strictly necessary for the production of an 'acceptable' target language word order. As we saw in section 6.4, most of these changes are determined by specific information attached to source language lexical items; very little is controlled by general operations. The inextricable mixture of transfer and generation is now seen as a major disadvantage of direct systems (particularly those of the 'first generation'). The lack of modularity means that improvements are almost impossible once the systems are in operation. Later direct systems have increased modularity, and some separation of 'transfer' and 'generation' procedures is found, although the basic orientation of the processing to specific source and target languages remains. Examples of systems which exhibit these features are given in the descriptions of Systran (Chapter 10) and of Météo (Chapter 12).

7.2 Generation in indirect systems

In general, indirect systems reflect a more sophisticated attitude to the linguistic and computational aspects of MT system design. Nevertheless, when it comes to target language generation, several systems still go about it in a rather crude manner. First, it is not infrequent to find a mixture of transfer and generation: they may well have a distinct and independent analysis phase (i.e. one that takes no account of target language structures), but they do not necessarily make a clear separation of transfer and generation procedures (for example, METAL, described in Chapter 15). Second, the basic approach in most cases can be characterized as 'pre-determined' in that only the first and most straightforward possibility is considered at each point in the generation process. The disadvantages of this approach is discussed later (section 7.3) after describing typical generation procedures.

7.2.1 Generation in transfer-based systems

In a transfer system, the generation phase is generally split into two modules, 'syntactic generation' and 'morphological generation'. In **syntactic generation** the intermediate representation which is output from analysis and transfer resembles a deep-structure tree of the older type of transformational-generative grammar (cf. example (27) in section 6.5). It is converted by 'transformational rules' into an ordered surface-structure tree, with appropriate labelling of the leaves with target language grammatical functions and features. The basic task of syntactic generation is to order constituents in the correct sequence for the target language.

For example, if a sentence is labelled 'passive' in the deep structure, syntactic generation will create a node for the auxiliary verb, labelled with the appropriate tense information, and will assign the 'past-participle' label to the main verb. Similarly, if a noun phrase has a clausal modifier and this has to appear as a relative clause in the target, syntactic generation involves the selection of an appropriate relative pronoun and an appropriately tensed verb form. Sometimes, an argument of a verb which was, for example, the direct object in the source language has to be a prepositional object in the target language (cf. French *chercher* and *regarder* with direct objects, and English *look for* and *look at*). Assuming that transfer retains the simpler direct object construction in the intermediate representation, it will be in the syntactic generation module that the prepositional phrase structure is created.

The resulting 'surface' structure is then input to **morphological generation**, which interprets strings of labelled lexical items for output as target sentences. It is generally relatively straightforward: the labelled string '*dog*+plural' to give *dogs* in output, German '*machen*+past+plural' to give *machten*, etc. In some cases morphological rules are a little more complex (e.g. French '*remplir*+imperfect+3rd-person+plural' → *remplissaient*), and of course there must be special rules to handle irregular forms ('*go*+past' → *went*). In general, morphological generation can usually be handled by a combination of general and special-case procedures, on a word-by-word basis.

It is appropriate at this point to look at a simple example based on Ariane (a system described in detail in Chapter 13). This example is derived from the German-to-French version, with some details altered for the sake of clarity. The representation for the input sentence (1) after lexical and structural transfer is given in (2).

(1) *Das Energieproblem ist zu einer grossen Herausforderung geworden.*
'The energy problem has become a major challenge'

It should be noticed that the basic structure of (2) is somewhat like that of the German sentence: the main verb is still at the end, and some of the values of label features also derive directly from the German analysis. For example, the analysis of *Herausforderung* ('challenge') as a derived verbal noun in German has been retained in the features for the French equivalent, given here by the underlying verb *défier*. Other features have been inserted during lexical and structural transfer, particularly the values for the lexical units themselves.

(2)

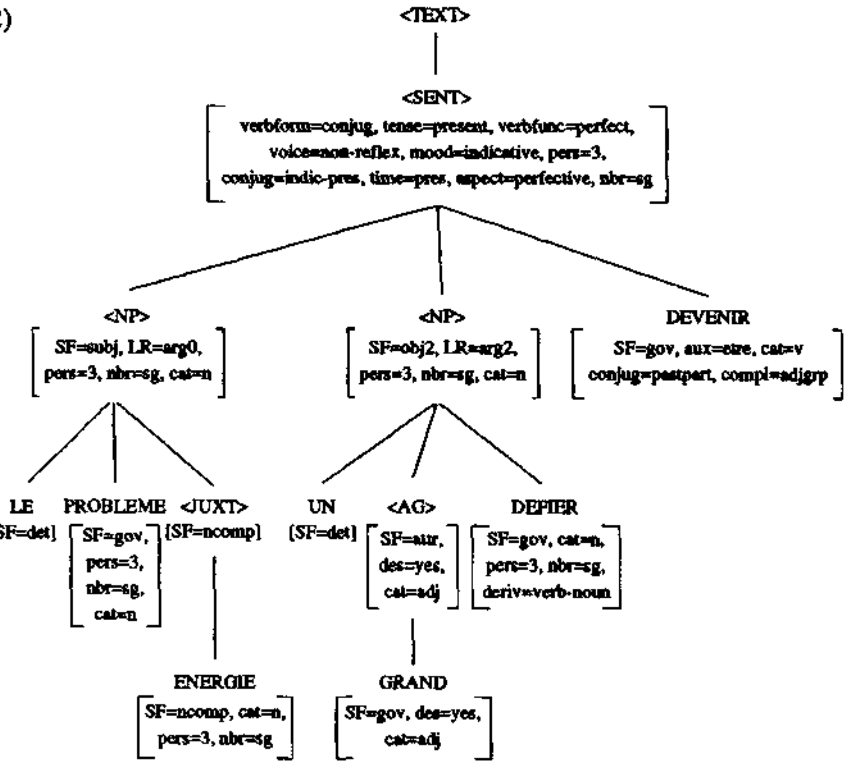

The task of syntactic generation is to convert this representation into a surface structure for the corresponding French sentence. It involves for example looking at the labelling on the top node <SENT> in order to select an appropriate verb form. In this case the labels 'perfective' and 'present' indicate the introduction of an auxiliary, and the 'aux=etre' label on the DEVENIR node specifies which auxiliary it should be. In a similar way, the labelling of <JUXT> as 'ncomp' indicates that the noun complement *énergie* must be transformed into a prepositional phrase introduced by *de*. A further task of syntactic generation is the distribution of number and gender information to relevant terminal nodes. For example, determiners and adjectives should agree in number and gender with the governing noun in a noun phrase; the auxiliary verb, which is to show the tense (from features on <SENT>), should agree in number with the subject noun (phrase); and the past participle should also agree in number and gender with the subject. The resulting surface structure tree (3) may still show constituents from the intermediate representation, but they are ignored in final output.

This structure is now subjected to a simple **morphological generation** routine. This works through the tree top-down and left-to-right. Each time it comes to a terminal node, the set of labellings are interpreted according to morphological rules. For example the node labelled LE [nbr=sg, gend=f, vwl=yes] corresponds to the string *l'*; ETRE [tense=pres, pers=3, nbr=sg] to *est*; GRAND [des=yes, nbr=sg, gend=m] to *grand*; and so on. The result is the surface string (4).

(3)

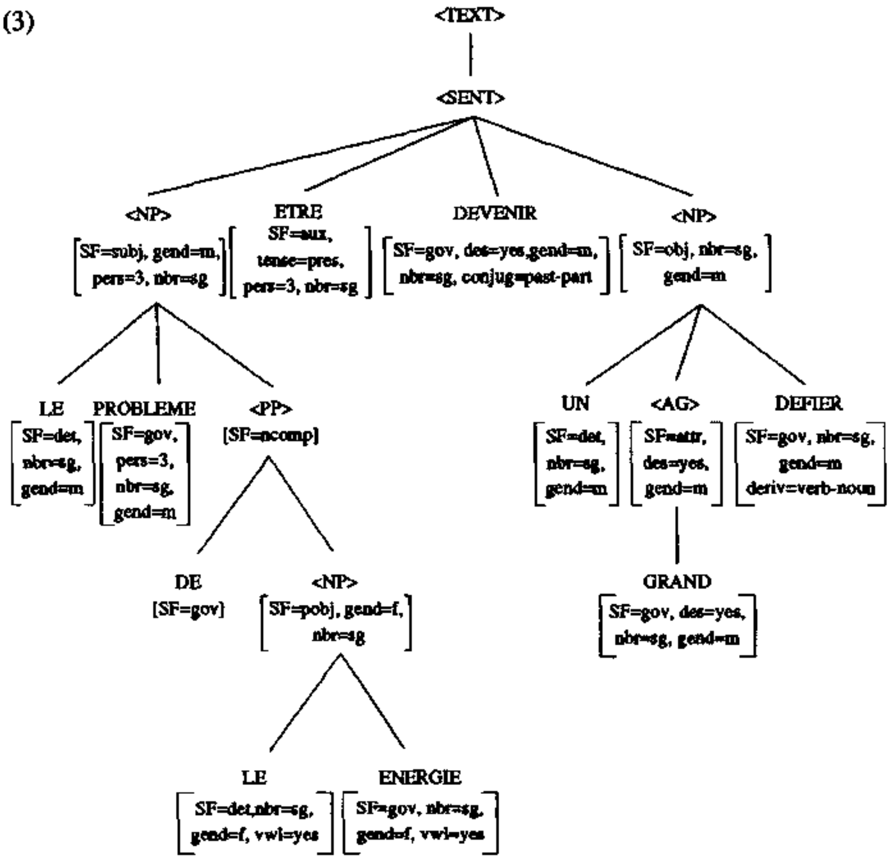

(4) *Le problème de l'énergie est devenu un grand défi.*

Morphological generation is generally no more complex than this: the problem of generating *l'* rather than *la* in (4) can be handled by a kind of agreement feature ('vwl=yes'), though it should be noted that such essentially phonetic rules are sensitive to word order rather than grammatical constituency. For example, the choice between *a* and *an* in English is determined by the next word in the sentence, not by the head noun. Other problems for morphological generation include the choice of different stems for different endings, e.g. *good+er* → *better*, French *aller*+2nd person+sing+present → *vais*. There are however a number of rather more subtle problems. Sometimes a morphologically derived form is available as an alternative to a full form: for example, *painter* and *artist* (which differ little semantically). And, just as in morphological analysis, compounding can be a major headache, when some means has to be found for deciding whether or not elements should be recombined as compounds, e.g. *drainage filtering system* vs. *system for filtering during drainage*, or German *System für Übersetzung* vs. *Übersetzungssystem* ('translation system').

7.2.2 Generation in interlingua systems

In essence, the procedures for generating texts in interlingua-based systems parallel those described for transfer-based systems. Generation includes stages of syntactic and morphological generation. The main difference is that the starting point, the input, is not a 'deep-structure' syntactic representation, such as (2), but an interlingua representation based probably on predicate–argument structures. Examples of such structures were given in section 6.7. The 'deep' syntactic structure must first be generated from the interlingual representation by a stage often known as **semantic generation**.

The process may be described using example (5), which has been adapted from (37) in section 6.7. As before, interlingual lexical items in the representation are indicated by angled brackets, with English labels for convenience.

(5)

Pred = <DEPEND>
Tense = timeless-truth
Patient = [Pred = <GOVERNMENT> ; Num = sing ; Qtf = universal]
Inst = [Pred = <SUPPORTER> ; Num = plur ; Mod = [Pred = Pron ; Ref = → (Patient)]]

The 'deep' structure to be generated is (6) (also given as (27) in section 6.5).

Starting with the topmost Pred <DEPEND>, the 'semantic generation' module must select a verb in the target language (in our case, French) as head of the sentence. The choice of *dépendre* entails the formation of a complement structure ('fun=compl') corresponding to the Inst in (5). This consists of an appropriate prepositional 'case' marker and a dependent noun phrase, corresponding to the structure of the Inst, with as its head the target language equivalent of <SUPPORTER>, i.e. *défenseur*, and the feature 'num=pl' derived in an obvious way from the 'Num=plur' feature in (5). At the same time, the feature 'timeless-truth' is attached to the top node (S) as the target language feature 'tns=pres'. The Patient feature structure in (5) is converted into an NP with the feature 'fun=subj', with the noun *gouvernement* (the target equivalent of <GOVERNMENT>) as its head, and its quantifier (Qtf=universal) giving the determiner *tout* and the singular number for the whole NP. Finally, the generation module produces a representation of the reference link (Ref) from the Mod of <SUPPORTER> to the subject NP. It is achieved by creating a possessive pronoun (pospron) with a referential co-index to the subject NP (i.e. attaching the feature 'refer=*1' to both nodes): for French (unlike English with *his*, *her*, *its* and *their*), only the number of the antecedent

(6)

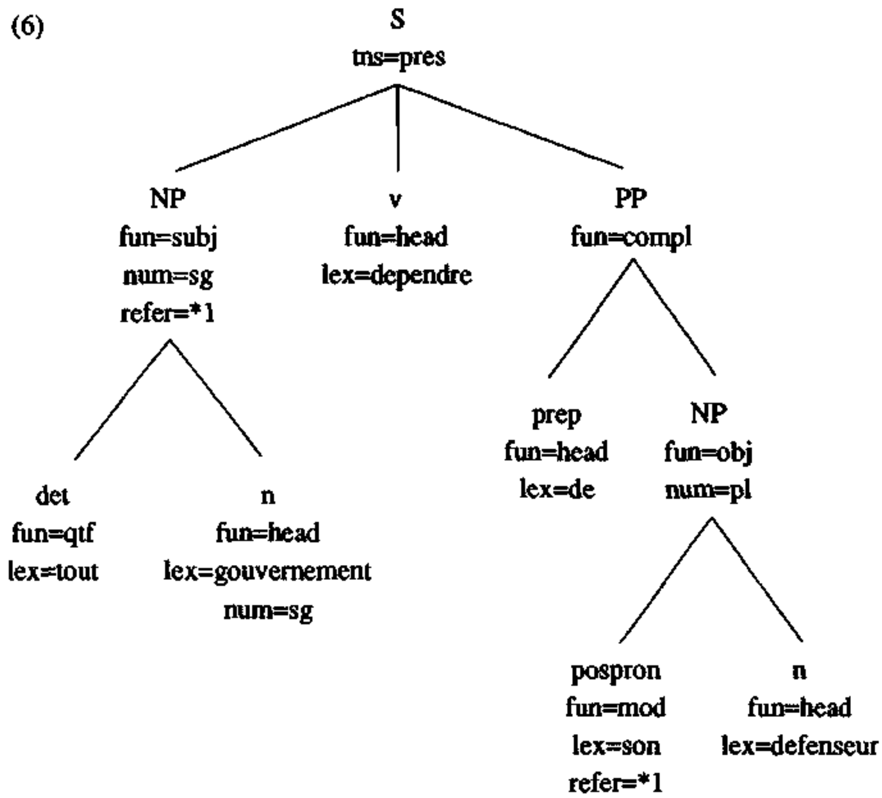

is important for choosing the appropriate possessive pronoun (*son* or *leur*, though number and gender agreement with the noun it modifies may result in a different surface form, e.g. *sa*, *leurs* etc.). The number information from the antecedent is added to the representation.

Subsequent stages of syntactic and morphological generation would then proceed as outlined above for a transfer-based system. The result would be (7a), the translation of (7b) (example (17) in section 6.3).

(7a) *Tout gouvernement dépend de ses défenseurs.*

(7b) Any government is dependent on its supporters.

7.3 Pre-determined choices and structure preservation

In certain respects the approaches to generation in indirect systems are merely more sophisticated versions of the approach in the earliest direct systems. There are two major problems with this model: the assumed 'pre-determination' of target language generation, and the implication of 'structure preservation' as a methodology in translation.

In the context of generation **pre-determination** means that for any given deep representation that is derived from an interlingua or from transfer, there should be one (and only one) corresponding surface structure. The opposite approach would be where choices have to be made between the possible surface structures which may be generated from the same representation.

The reasons for adopting pre-determination as a strategy in generation seem to be computational: a generation module in which the correspondence between one representation and another (deep and surface syntax, or intermediate representation and text in the extreme case) is easier to implement and maintain. It means that once transfer is complete, the remainder of the translation process is a straightforward process. It may involve a lot of computation, but it should not be complicated. The alternative entails the development of algorithms for choosing appropriate structures; it introduces added complexity.

The assumption of pre-determination goes against what is known both by linguists and by translators: that there is no one-to-one correspondence between meaning and form. From the linguist's point of view, this is expressed by observing that, just as one surface string may have several different possible deep structures (structural ambiguity, section 5.3 above), so a given deep structure can be mapped onto different surface structures. In Chomskyan linguistics this fact was captured by the postulation of optional transformations which preserved meaning (section 2.9.2) A translator knows very well that there are hundreds of ways that a given text can be expressed in the target language: in fact, most translators would regard the task of expressing the meaning of a source text, rather than initial understanding of the text, as the aspect of translation which requires most skill.

Although MT is not, of course, the same as human translation, the somewhat naïve assumption that generation should be pre-determined in this sense has the result that systems tend to produce rather stilted translations with very little variety in style. In itself, this may be seen as an advantage, especially when MT is used for technical translations, where the main requirement is accuracy and fidelity to the original (see Chapter 9); but on the whole the principle results in MT output of a lower quality than it might be.

The basic problems are too close adherence to the structures of source texts and the treatment of sentences independently. To avoid making choices during generation, it is seen as preferable to retain the structure of original sentences in texts: **structure preservation** is adopted as the option of first choice. If at all feasible, the structure of the source language is preserved in the generation of the target language; and, as a result, translations are as literal as possible. Rearrangement of structures takes place only under certain predefined conditions. In this respect, MT is in direct contrast with human translation, where preservation of sentence structure is normally the LAST choice; even the sequence of sentences is often changed. Only when the exact meaning of the source text is perhaps unclear (and the original author is unavailable for clarification), do translators generally resort to translating more or less phrase by phrase.

Text generation is a major concern of many researchers outside MT in the field of natural language communication, in particular those interested in database interrogation. In this context the aim is to respond in natural language to users' requests for information. The factual content of the texts to be generated is derived from the internal representation of the information in the database. Initially, the most popular technique for generating responses was that of **direct replacement**, in which representations were merely transformed into natural language sentences more or less in the pre-determined manner described above. The inflexibility of the approach became soon obvious. It was clear that the most appropriate and natural reply to a query ought to take into account the context in which it was made; it should not make unnecessary and unnatural repetitions, and it should relate to what has gone before, i.e. the generation of texts has to make full use of ellipsis, pronominalization, and topic-comment structuring, any of which might differ from one language to another (consider how in examples (5) and (6) above it was assumed that the possessive pronoun in the source text should appear as a possessive pronoun in the target.) Text generation is now seen as a series of decision points involving choices conditioned by a vast range of information including, but by no means restricted to, the propositional content of the texts to be generated.

This approach to text generation would take into account information about the pragmatic situation, whether 'new' or 'old' information is being communicated, what links are being made to previous sentences, etc. (section 2.7) For example, consider again the sentences in (8) (which we saw in the last chapter).

(8a) *Der Mann jagte den Haifisch.*
THE MAN-subj CHASED THE SHARK-obj

(8b) *Den Mann jagte der Haifisch.*
THE MAN-obj CHASED THE SHARK-subj

If the input representation to the generation module were to retain the 'subject' specification of the original, i.e. *Mann* in (8a) and *Haifisch* in (8b); and if the rules for generating English state that the 'grammatical subject' should appear first in the sentence (which is the norm in English), then the resulting translation would be as in (9).

(9a) The man chased the shark.

(9b) The shark chased the man.

But this would be wrong for (9b). While (8a) and (9a) have both the same topic (*Mann* and *man*), (8b) and (9b) do not (*Mann* and *shark*). In fact, (9b) corresponds to (10a). What should happen is that the marked thematization of (8b) should be preserved (cf. section 2.7). In order to retain the same topic-comment structure, *man* should be the first noun in the English. This means either a 'focus' construction (10b) or a passive structure (10c).

(10a) *Der Haifisch jagte den Mann.*
THE SHARK-subj CHASED THE MAN-obj

(10b) It was the man that the shark chased.

(10c) The man was chased by the shark.

Other factors which should influence text generation come broadly under the heading of 'style'; for example, the choice between an adjectival phrase and a relative clause (11), between different types of passive (12), and between alternative expressions for generic nouns (13).

(11a) the late trains

(11b) the trains which are late

(12a) The cake is eaten.

(12b) The cake gets eaten.

(12c) Someone eats the cake.

(13a) The lion is dangerous.

(13b) A lion is dangerous.

(13b) Lions are dangerous.

In the context of translation two further points should be made. First, the kinds of choices just exemplified are often conditioned by quite different criteria in different languages: for example, heavy pre-nominal modification is quite permissible in German but only rarely in English, cf. (14) and (15).

(14) *die in den letzten zehn Jahren entwickelnden Industrien*
THE IN THE LAST TEN YEARS DEVELOPING INDUSTRIES
'the industries (which have been) developing in the last ten years'

(15) the above mentioned problem

Second, the syntactic devices used to express various choices are not always equivalent between languages: German expresses thematization by word order (16a) while this is not always possible in English (16b) or French (16c). English has available a variety of alternative devices (16d,e). In these examples the issue of quantifier scope is also involved, a problem in generation just as in analysis (cf. section 5.5).

(16a) *Nur Kuchen hat das Kind gegessen.*

(16b) * Only cakes has eaten the child.

(16c) * *Seulement des gâteaux a mangé l'enfant.*

(16d) The child has eaten nothing but cakes.

(16e) It is only cakes that the child has eaten.

7.4 Stylistic improvements in generation

The structure-preserving approach continues to be popular for good practical reasons. On the one hand, problems of analysis and transfer are complex and difficult to resolve. If a simple version of generation produces reasonable results — ones which, though not perfect, can be post-edited relatively easily or serve the purpose of conveying more or less accurately the meaning of the original text to the reader — then it makes sense to adopt it. On the other hand, until fairly recently most MT systems have translated between Indo-European languages with broadly similar basic structures. As a general rule, for these languages a verb in the source language corresponds to a verb in the target language, an adjective corresponds to an adjective, a subject–predicate structure remains more or less constant, and so on. Exceptions can be treated by special procedures in 'structural transfer' (section 6.5.2). But the basic assumption can be maintained that it is in fact possible to retain source language sentence structures.

The inadequacy of this approach became obvious when more MT research tackled translation between languages of greater typological diversity (notably English and Japanese). One example was illustrated in the previous chapter: the so-called 'animate subject constraint' in Japanese. An even greater difference between the two languages is the 'action-orientation' of English and the 'process-orientation' of Japanese. In English it is more natural to express events in terms of someone doing something, while in Japanese it is the tendency to focus on the object undergoing the action. For example, a natural translation of (17a) would be (17b) rather than (17c). Although grammatical and meaningful, sentence (17c) refers to the physical capacity of sight and might imply that the speaker has been cured of blindness.

(17a) I can see Mount Fuji.

(17b) *Fuji-san ga mieru.*
'Mt. Fuji can be seen'

(17c) *(Watashi wa) Fuji-san wo miru koto ga dekiru.*
'It is possible (for me) to see Mt. Fuji'

In confronting these problems, a clear division between the respective roles of analysis, transfer and generation needs to be maintained. The analysis module should be expected to deal with source language ambiguities, the transfer module with questions of contrastive lexical and structural differences, and the generation module should be expected to handle any problems which are matters of choice in the target language. In all the examples so far the criteria for choosing among different expressions of a given 'meaning' can be stated without reference to the source language: the differences between the alternatives in examples (11)–(13) are features of English, and are quite independent of the way the same ideas might have been expressed in a source language, be it French, Russian or Japanese.

The first requirement, therefore, is some means in the generation module for determining the criteria to be applied in particular instances. Sometimes this is easier than others: a structure-preserving translation into English of (14) would not actually be completely ungrammatical, but in order to block its generation there would need to be a 'preference rule' stating that noun modifiers consisting of anything much more than an adverb and an adjective are not to be produced. Similarly it should be possible to formulate a rule which would allow (18a) but prevent a large noun phrase appearing as an indirect object before a short direct object as in (18b).

(18a) John gave the man a book.

(18b) * John gave the man he met yesterday in the park after he'd been for a hamburger a book.

On the other hand, without considerable contextual and pragmatic information it might be impossible to choose between the two variants in (19).

(19a) The boy gave his sister a book.

(19b) The boy gave a book to his sister.

More complex and problematic would be the coding for choosing from amongst the permutations of translations for the German sentence (20a), where for each component part in the English representation — very roughly as in (20b) — there are various possibilities (20c).

(20a) *Technologische Fortschritte sind ein entscheidender Bestandteil der Entwicklung, insbesonders für Länder, die für ihr Energiebedarf vom Import abhängen.*
TECHNOLOGICAL ADVANCES ARE A(N) ESSENTIAL INGREDIENT OF-THE DEVELOPMENT ESPECIALLY FOR COUNTRIES WHICH FOR THEIR ENERGY-NEED FROM-THE IMPORT DEPEND

(20b)

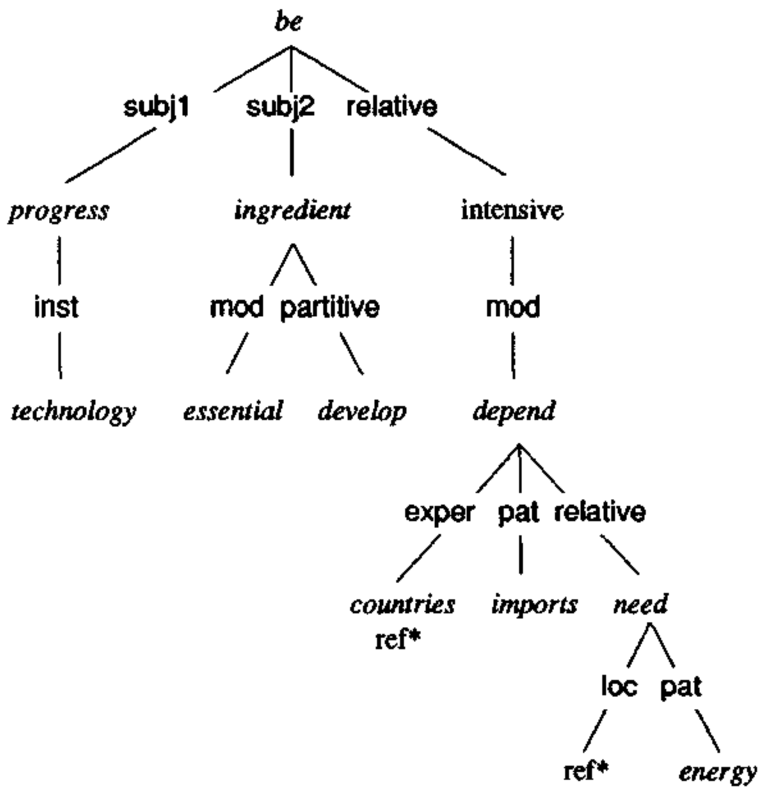

(20c) Technological progress/advance(s)...
Progress/Advance(s) in technology...
... is/are an essential ingredient of development...
... is/are an ingredient which is essential for/in development...
An essential ingredient of development is technological progress/advance(s)...
... especially/particularly for (those) countries...
... (which are) dependent on imports for their energy needs...
... which depend on imports for their energy needs...
... whose energy needs are dependent on imports...
... whose energy needs depend on imports

It should be noticed in particular that whatever the criteria for making the choices are, the corresponding choices in the German source text are more or less irrelevant.

An important point to note about this view of generation is that it requires a high degree of language independence, i.e. that intermediate representations should go further towards neutralizing the accidental syntactic characteristics of source languages. It is a problem for both transfer-based and interlingua systems, but in this case the motivation is clear: the interface structure should contain just the information that is necessary to permit appropriate decisions to be made by the generation module.

In this section, we have suggested that there are problems of generation which are just as complex as, though subtly different from, those of analysis, and that in most MT systems they have not been properly addressed. Generation should not only handle the mundane questions of morphological agreement and basic word order, but should also tackle what for a human translator might in the end be classified as matters of style. Some help in the selection of appropriate output may well come from access to databases of 'example' translations (as described in section 6.9), but more research is needed on computational stylistics, and particularly computational implementations of the findings of comparative stylistics. The need is for methods of ensuring the generation of structures which are idiomatic for the target language, e.g. the 'process orientation' of Japanese, the 'action orientation' of English in (17) above, and the preference in French for more 'dynamic' expressions (21a) than in English (21b).

(21a) *Les gens ont applaudi sur le passage des troupes.*

(21b) People cheered as the troops marched by.

The task is far from easy, but it must be tackled if significant improvements in MT output are to be achieved.

7.5 Generation as the 'reverse' of analysis

It is a common assumption by newcomers to computational linguistics in general and to MT in particular that generation ought to be the mirror image of analysis, and the assumption is probably reinforced by diagrams of indirect MT systems, by diagrams such as Figure 6.1 (Chapter 6) and by some of the general descriptions of systems. Our discussion so far, in this and previous chapters, especially on analysis procedures, should have dispelled this misconception.

Nevertheless, there has recently emerged a style of programming based on logic, using the Prolog programming language, in which this literal interpretation of generation as the reverse of analysis can actually be realised. Prolog is claimed to be a non-procedural programming language, which means that programs written in it do not specify how something is to be computed but describe what has to be computed. Prolog programs describe sets of data and logical rules. The sequence in which the former interact by the application of the latter is determined by an 'interpreter', which decides which rules are appropriate, and when to apply them. Only when there are several candidate rules that could be applied does the order in which the programmer happened to write them down have any significance.

In this respect Prolog is said to embody the 'separation of algorithms and data' discussed in Chapter 3 (the data are the programs; the algorithms are more or less hidden from the programmer), and this is hardly surprising since, in fact, the development of Prolog grew out of work on Q-systems in what later became the Météo project (see Chapter 12).

One consequence of this program design is that Prolog programs — at least simple ones — are, under normal circumstances, reversible. A Prolog program defines a set of relationships between sets of data: in the case of a parser, the program will define a relationship between a set of possible texts (strings of words) and the corresponding representations. Because of the way Prolog works, we can in principle use EXACTLY THE SAME PROGRAM either to compute a representation given a string of words or to compute the string of words given the representation. In this sense, a Prolog program written as a parser can also be used as a generator: we can use it to ask either "What is the structure corresponding to this sentence?" or to ask "What is the sentence corresponding to this structure?"

A number of researchers have adopted this basic idea in experimental MT systems and are investigating the limits of its applicability. Using Prolog they are attempting to build reversible (typically bilingual) MT systems where the grammars used for analysis are also used for generation just by running them backwards. Unfortunately it is often the case that in striving to maintain reversibility the grammars remain too simplistic, and the system ends up as a low-level interlingual one where the intermediate representation is little more than a surface phrase-structure tree: only sentences where there are no gross structural differences (e.g. of the *like/gern* type) can be translated. When structural changes involve deleting parts of the structure, reversibility becomes much more problematic: a reversed deletion is an insertion, and it is often difficult to control the arbitrary insertion of underspecified parts of a structure.

An alternative to this 'reversible procedural' approach involves the use of feature unification (see section 2.9.7) and the definition not of reversible analysis grammars as such, but of collections of well-formedness conditions or constraints on feature bundles, which, when applied to one structure (which might be a simple string of words) combine and unify to give a new structure. The approach can be used both for analysis, where a string is given as input and a representation 'emerges', or for generation, where the representation is given as input. In order to translate, the constraints must also specify mappings between representations, so that input can be the source string and output the target string. The Eurotra 'E-framework' (Chapter 14) is in part an implementation of this idea, though for practical reasons there are aspects of the formalism which make it not fully reversible.

It is frequently assumed that reversible grammars are most appropriate for interlingua systems, where the 'analysis' and 'generation' modules share a common interface; but there is nothing against a transfer-based reversible system, as long as the transfer rules themselves are also reversible.

In both types of 'reversible' system, an additional issue is the general underlying principle that translation itself is in some sense reversible, i.e. that because a sentence S can be translated as S′, then S′ must be translated as S in the opposite direction. Unfortunately, this is far from the case in reality. Nevertheless, several basic research MT systems have taken the challenge of reversibility seriously, since there is a clear reduction in development costs to be had if large amounts of linguistic data can be used for both analysis and generation.

The best known example of an MT project explicitly investigating reversible grammars is Rosetta (described in Chapter 16). This interlingua model makes the further theoretical claim that grammars should not only be reversible, but that they should be 'isomorphic', i.e. that for every rule needed for one language there should be a corresponding rule in the other language. Other reversible MT systems written in Prolog make less strong theoretical claims. They include ULTRA, a multilingual system with all combinations of English, German, Spanish, Japanese and Chinese, under development at New Mexico State University, Las Cruces; a related project XTRA, investigated at various centres as its developer changed laboratories; SWETRA at the University of Lund, developed chiefly for English–Swedish, but with parallel experiments for French, Polish, Russian, Georgian and Irish translation; and CRITTER, a sublanguage system developed in Canada to translate agricultural market reports between English and French (see section 18.7) In each case, analysis and generation are both handled by definite clause grammars (DCGs), a formalism for augmenting essentially phrase structure rules, developed specially for Prolog programming.

7.6 Further reading

General discussions of the problem of generation are very scarce in the MT literature. For example, the section of Nagao's (1989) book on 'sentence generation' is actually about differences in English and Japanese sentence structure, and so is really about transfer, while McDonald's (1987) article is a view of MT generation from an 'outsider' working on information retrieval systems. We must consult descriptions of the target-language generation modules of individual systems; the example of Ariane generation in section 7.2.1. comes from Guilbaud's (1987) description of the German–French version.

One notable research project which was chiefly interested in generation was the SEMSYN project (Laubsch *et al.* 1984, Emele *et al.* 1986, Rösner 1986a,b) in which a German generation module was devised for producing output from the intermediate representations of a Japanese–English MT system.

The issue of pre-determined choice in target-language generation is hinted at in various descriptions of the Ariane system (for references see Chapter 13), and the fact that the correspondence between a text and its possible translations is one-to-many rather than one-to-one is implicit in several early discussions of transfer in Eurotra and the notion of 'isoduidy', an invented term indicating the relationship between two translation equivalents (or more generally, linguistic representations which are equivalent alternatives at a given level of abstraction) (see Johnson *et al.* 1985, Arnold 1986).

The problem of structure preservation as a first choice is discussed in Somers (1987b, 1990) and in Somers *et al.* (1988).

An important contribution to the discussion of 'stylistic' improvement of MT has been made by DiMarco and Hirst (1988, 1990).

The example with *Technologische Fortschritte* (in section 7.4) and a discussion of the requirements for intermediate representations come from Somers (1987b).

For more information on ULTRA, see Farwell and Wilks (1991). The XTRA system is reported in Huang (1988). See Sigurd (1988) for a discussion of SWETRA. References on Rosetta and CRITTER are given in Chapters 16 and 18 respectively. For information on the use of Prolog for natural language processing, and on DCGs, textbooks such as Pereira and Shieber (1987), Gazdar and Mellish (1989) or Gal *et al.* (1991) are recommended.