The Re-use of Linguistic Resources across Languages in Multilingual Generation Components

John Bateman
GMD/Institut fur Integrierte
Publications- und Informationssysteme
Dolivostrafic 15, D-W-6100 Darmstadt, Germany
USC/Information Sciences Institute
4676 Admiralty Way, Marina del Rey, CA 90292, USA

Christian Matthiessen, Keizo Nanri, Licheng Zeng
Department of Linguistics
University of Sydney
Sydney, N.S.W. 2006
Australia

Abstract

An approach to generation system design is described which supports maximal expression of commonalities across languages. Within this approach it becomes natural to represent inherently multilingual grammars and semantics. The approach rests on the linguistic notion of *functional* similarity and difference: by capturing the functions languages need to perform, we achieve a level of linguistic description which carries across languages far more effectively than accounts that are structurally based. We demonstrate the general principles, implementation, and benefits of the approach with respect to three unrelated languages: English, Chinese, and Japanese.

1 The task of multi-language generation

Despite the growing need to develop text generation systems capable of producing texts from the same knowledge source in more than one language, most major generation systems remain mono-lingual: there has also been very little work on extensible multi-language generation that seeks an architecture within which the work involved in adding a new language may be minimized. Reliance upon an appropriate typological theory of language, however, makes it possible to build a system where different languages can both converge and diverge in terms of knowledge source, semantics, grammar, etc. In this paper, we describe such an approach to generation system design, showing how it supports maximal expression of commonalities across languages. Within this approach it becomes natural to represent inherently multilingual grammars and semantics that provide a basis for the rapid development of new grammar and semantic components. The approach rests on the linguistic notion of *fundtional a* /similarity and difference: by capturing first the functions that languages need to perform, we achieve a level of linguistic description which carries across languages far more effectively than accounts that are structurally based.

The paper demonstrates the general principles, benefits, and implementation of the approach we propose with respect to three languages that are unrelated according to structural typologies and language families: English, Chinese, and Japanese. We will see that despite the superficial differences between them, the communicative tasks that all three need to achieve in fact enforce a degree of commonality of description that can be used to support maximal sharing of linguistic resources.

2 Linguistic description based on function: a monolingual example

In order to build a theoretically sound system based on a general, typologically informed account of multilingualism, we have begun a collaborative research effort that takes a well-established English text generation system, the Penman system (under development at USC/ISI since 1980), as its base. This choice is motivated by the *fun*ctional-orientation of linguistic description within the system. Our basic assumption is that commonality across languages is *functional* in the first instance, not structural or realizational: i.e., functionality has to be preserved across languages, but structural realizations may very well differ. To begin our description of the multilingual system, therefore, we first illustrate the kind of linguistic description which provides a basis for multilingual!ty: i.e., the functionally oriented description of the Nigel systemic-functional grammar [Matthiessen, 1985] for English developed within the Penman system. Systemic-functional grammars have the required property of representing function as theoretically prior to structure and this, as we shall see, is the crucial perspective which makes multilingual descriptions practicable.

A systemic-functional grammar captures grammatical variation in terms of *minimal grammatical choices* between grammatical features. These choice points are organized into an interconnected network representing the interdependencies between possible grammatical feature selections. The grammar of a language is the complete statement of these interdependencies and their consequences for structural realization: this is the *systemic network* from which the general theory takes its name. Crucially, the principle motivations for including features are *functional* rather than structural. Grammatical structures are only arrived at by making these abstract functionally motivated choices of features in the network: each feature selected can then add structural constraints to the specification of the linguistic unit being described.

More formally, a systemic network representing functional organization is a directed acyclic graph with labelled arcs whose nodes correspond to grammatical choice points; these are called *systems.* The outward directed labelled arcs denote the *terms of* the system. Each system has two or more terras, or output features, which at the stratum of grammar represent minimal grammatical alternations. In addition, the inward directed arcs for each system denote an *entry condition* which determines the paradigmatic context in which the alternation represented by the system is relevant. In general, the entry condition is defined as a boolean expression over the possible input are labels.

As a particular example area, we consider the resources of MOOD in English grammar: that is, the area of the grammar concerned with speech act: questioning, asserting, etc. Figure 1 shows a simplified extract from the MOOD fragment of Nigel. Here we can see that there are three grammatical systems MOOD TYPE, INDICATIVE TYPE, and INTERROGATIVE TYPE. The structural consequences of the grammatical features are shown in boxes underneath the feature which controls their application. Thus, we can gloss this network fragment as follows. The speaker must make an abstract choice between whether s/he wishes to make an indicative clause or an imperative clause; if an indicative clause is chosen, then several structural constraints apply, i.e.: the insertion of functionally labelled constituents Mood, Subject, and Finite, and the structural positioning of Subject and Finite as subconstituents of Mood. Then, the speaker must go on to make an abstract functional choice between 'declarative' and 'interrogative'; if a declarative clause is selected, then the structural constraint that the constituent labelled as the Subject must precede the constituent labelled as the Finite applies. Finally, iff the speaker had instead selected 'interrogative', then s/he must also go on to select between 'polarity' and 'element'[1]; 'polarity' calls for the structural constraint that the Finite precede the Subject (giving rise to yes/no questions: *did he go?),* 'element' calls for the structural constraint that there is a functional element labelled Wh, and that this element occurs first — i.e., is conflated with the constituent Theme which always comes first in English clauses (giving rise to Wh-questions: *where did he go?).*[1]

Each abstract choice point in the grammar network also has a decision procedure associated with it that decomposes the minimal grammatical alternation into the minimal semantic distinctions necessary for motivating that alternation. The decision procedures are called *choosers* and the individual semantic alternations are called *inquiries* [Mann, 1983]. Although our account of multilinguality extends into these areas also, space precludes their discussion here. We shall for present purposes simply take the level of inquiries as given and note

[1]For a more complete account of the Nigel grammar of English and the linguistic motivations for its descriptions, see [Halliday, 1985, Matthiessen, 1990]. For an almost complete account of the use to date of systemic-functional linguistics in (monolingual) text generation, see [Matthiessen and Bateman, 199l).

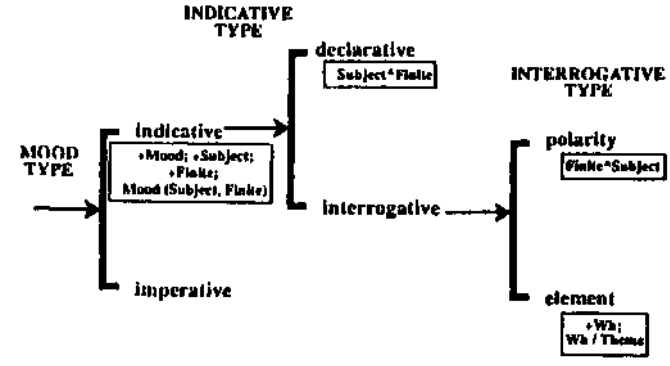

Figure 1: Grammar fragment for English speech acts

that the set of semantic distinctions for the MOOD systems are: command-q, statement-q, question-q, and polarity-variable-q; these may for the purposes of this paper be seen as classifying speech acts sernantically in the obvious way suggested by their names wr see something of their use below in Section 4.

3 The move to multilinguality: Chinese and Japanese

Although the Nigel treatment of MOOD is part of a grammar of English, its functional orientation is not,, in fact, particularly English-specific. It seems very likely that all languages have evolved grammatical resources for establishing and maintaining symbolic interaction — i.e., for engaging in dialog and this basic dialogic resource can be understood more generally in terms *of* the metaphor of exchange [Halliday, 1984].

Given this level of abstraction, it is not surprising that many other languages also make use of this functional organization. In fact, we can up to a point use *identical* system networks fragments for Chinese and Japanese also. All three languages contrast, for example, declarative clauses with polarity interrogative ones. It is only when we come to more specific, or *delicate,* areas of functional potential that the languages begin to differ: whereas English stops in delicacy at polarity interrogative clauses, Chinese and Japanese provide further options. For Chinese, the speaker's expectation as to the answer to the polarity interrogative may be neutral ('tag' interrogative) or neutral/biased ('particle' interrogative) [Li and Thompson, 1981, Section 18.6]; and the interrogative tag may be either clause-medial or clause-final. Thus, in addition to simple polarity interrogative[4] just, Chinese, for example, provides the following further, more delicate functional distinctions:[2]

[2]There are, of course, further subtypes and distinctions: this is the case for almost all of the examples we give in this paper. Our purpose here is, however, to make the general principles and mechanisms clear rather than to provide exhaustive accounts.

tag, medial

ta wancheng	*mei wancheng*	*zuoye?*
s/he finish	not finish	homework

tag, final

ta wancheng	*zuoye*	*mei wancheng* ?
s/he finish	homework	not finish

Similar types of extensions in options are provided by Japanese. Furthermore, Japanese offers increased delicacy beyond the features 'declarative' and 'element'. The functions of these areas are as follows. In addition to giving information ('declarative'), a speaker of Japanese must also indicate whether s/he is the original source of that information ('originating') or not ('non-originating[1]): i.e., if the information was obtained from some other source, e.g., heard on the radio or generally rumored, then this must be overtly indicated. This is done by adding the expression *sou-desu/da* following the main process of the clause. The latter component of this expression is the copula verb 'to be' and it is this that carries the 'interpersonal deictic center' (Idc) of the clause: i.e., the part of the clause that carries statements concerning the social distance between the interlocutors (near: *da;* not near: *desu).* Also, in addition to demanding information concerning a particular element ('element'), a speaker of Japanese must also either leave the possibilities for that element unrestricted ('open'), which is equivalent to a simple Wh-question in English, or restrict the possible fillers of that element to a determinate set ('restricted'):

sore-	*wa*	*hon*	*desu*	*ka*	*nooto*	*desu*	*ka ?*
that-	TOPIC	book	is	QUESTION	notebook	is	QUESTION

"is that a book or a notebook?"

The increased functionality beyond polarity is similar for both Chinese and Japanese, although the more delicate variance between final or medial tags is not available for Japanese. Given the areas of convergence and divergence of functionalities we have now seen, it is possible to present a combined, multilingual grammatical system network for MOOD that covers English, Chinese, and Japanese; this is shown in Figure 2. Here we can clearly see the sharing of areas of grammatical description.

While English, Chinese, and Japanese thus have a number of MOOD systems in common, the various options in these systems are tactically different: that is, they are realized in different ways in the three languages. For example, in English the contrast between a 'declarative' and a 'polarity'-interrogative is realized by the Mood element, that is, by the combination of Subject and Finite; specifically, 'declarative' is realized by 'Subject before Finite', as in *he has,* and 'polarity' is realized by 'Finite before Subject', as in *has he.* However, Chinese does not have a Mood element consisting of Subject and Finite; the contrast between 'declarative' and 'polarity' is thus not realized by variation within such an element. Rather, the contrast is realized by the presence or absence of a special mood tag or particle; specifically, it is the 'polarity' interrogative that is marked by either a mood tag or a mood particle, depending on the type of 'polarity' interrogative, whereas in 'declarative', there is no presence of mood tag or mood particle. Japanese acts

```
(system :name INDICATIVE-TYPE
        :inputs indicative
        :outputs ((declarative
                    (order subject finite))
                  (interrogative))
        :chooser INDICATIVE-TYPE-CHOOSER)
```

Figure 4: Internal monolingual definition of a grammatical system in Penman

in this respect very similarly to Chinese; it does not have a Mood element and also realizes the contrast between 'declarative' and 'polarity' (in fact all 'interrogatives') by the deployment of a mood particle *(ka).* Now, since the grammatical theory separates the functional, strategic part of the grammar from the realizational, tactic aspect, stating these three sets of realizations separately is quite straightforward.

The realization statements for the three languages are tabulated in Figure 3; an example of their use is given in Section 4

4 Computational implementation of multilingual systemic networks

In this section we briefly describe the computational implementation of the multilingual environment for grammar description that we rely on to express generalizations of the kind illustrated in the previous section. To add the ability to define multilingual grammar components, we have extended the monolingual grammar definition notation provided by the Penman system. Figure 4 shows the current internal monolingual definition of the INDICATIVE TYPE grammatical system used above: the definition simply linearizes the information given graphically in the network notation; entry conditions for the system are given as a boolean expression over grammatical features (in this case simply 'indicative') under the slot :inputs and the output terms of the system, i.e., the grammatical features that are at issue for choice, along with their respective structural realizations, are given in a list under the slot :outputs. The value of the :chooser slot determines where to locate the decision procedure (or chooser) responsible for selecting purposefully between the grammatical features of the system in the context of a determinate text need.

We extend this definitional notation for system networks by allowing conditionalization according to language of all or some of the inputs, outputs, realization statements, and choosers. Language conditions are given as disjunctions *over* the languages for which the grammatical system definition is to hold specified by the keyword :variety. These disjunctions may be nested and cumulatively restrict the language applicability of the conditionalized components. We also allow conditionalization of entire systems by introducing a :variety slot for the definition as a whole,[3]

[3]We use the term :variety rather than, e.g., :language, to draw attention to language variety both across languages and *within individual languages.* Functional variety, or reg-

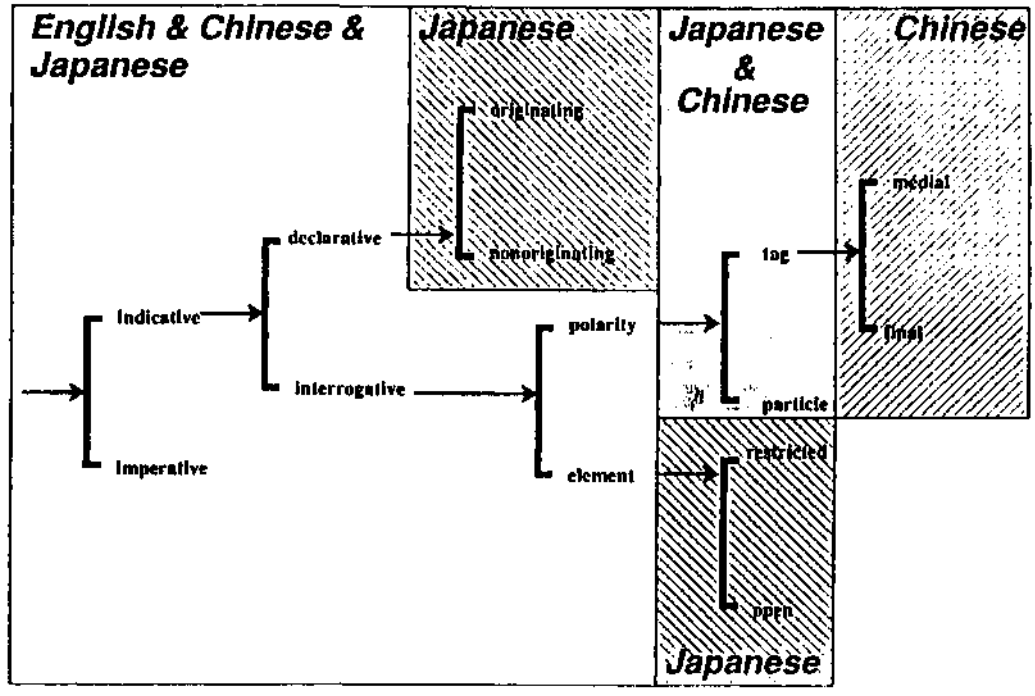

Figure 2: Multilingual MOOD network for English, Chinese, and Japanese

feature	*structural realizations*		
	English	Chinese	Japanese
indicative	+Mood; +Subject; +Finite; Mood (Subject, Mood)	–	–
declarative	Subject ^ Finite	–	–
interrogative	–	–	+Particle; Particle = *ka* Idc ^ Particle Idc / Process
element	+Wh; Wh / Theme	+Wh	+Wh
polarity	Finite ^ Subject	–	Particle ^ #
tag		+Tag; ...	Process (Positive, Negative)
tag & future aspect		Negator = *bu*	
tag & past/ perfective/ continuative		Negator = *mei*	
medial		Predicator ^ Tag	
final		Tag ^ #	
particle		+Interrogator; Interrogator = *ma*; Interrogator ^ #	–
originating			Idc / Process
nonoriginating			+Nonorgtag; Nonorgtag: copula; Process ^ *sou* ^ Nonorgtag; Nonorgtag / Idc
restricted			+Wh2; +Idc2; +Particle2; Idc2 = Idc; Particle2 = *ka* ; Particle ^ Wh2 ^ Idc2 ^ Particle2 #
open			Particle ^ #

^ : order
/ : unification
+ : constituent insertion
= : lexical assignment
: clause boundary
F (G, H) : subconstituency

Figure 3: Multilingual realization options for speech function

```
(system
 :name INDICATIVE-TYPE
 :inputs  indicative
 :outputs
   ((declarative
      (:variety English (order subject finite)))
    (interrogative
      (:variety Japanese (insert Particle)
                         (lexify Particle ka)
                         (order Idc Particle)
                         (conflate Idc Process))))
 :chooser INDICATIVE-TYPE-CHOOSER)

(system
 :name POLARITY-TYPE
 :variety (Chinese Japanese)
 :inputs  polarity
 :outputs
   ((tag
      (:variety Chinese (insert Tag))
      (:variety Japanese
          (expand Process Positive Negative) ...))
    (particle
      (:variety Japanese (order-at-end Particle))))
 :chooser POLARITY-TYPE-CHOOSER)
```

Figure 5: Example multilingual MOOD systems

```
(t / finish
   :actor (p / person)
   :actee (a / assignment)

        :command-q noncommand
        :statement-q nonstatement
        :question-q question
        :polarity-variable-q polarity-variable
        :expectation-q neutral
   )
```

Figure 6: Partial input specification (SPL)

Figure 5 shows example multilingual internal definitions required for some of our Chinese, Japanese and English MOOD systems given above. The first system, INDICATIVE-TYPE, is applicable to all languages covered, but only has structural consequences in English and Japanese. The second applies only to Chinese and Japanese. This means that all languages will have grammatical features such as 'interrogative' and declarative', but only Chinese and Japanese have the features 'tag' and 'particle'.[4]

Given this implementation, we can now generate clauses in multiple languages given a single knowledge base specification — as long as the information required by the distinct languages is common or more specific information than necessary is available. When divergent information is required in distinct languages, as, for example, occurs in our English, Chinese, and Japanese experiment when one language makes more delicate distinctions than another, then further inferencing may be required in order to obtain sufficient information for guiding generation. This can be seen more concretely in the following, where we provide an input specification that spans our three example languages. The input specification is the standard one used by Penman — the Sentence Plan Language (sPL: [Kasper, 1989]).

The partial SPL input specification[5] shown in Figure 6 states that there is a process t of semantic type finish with an actor p of semantic type person and an actec a of semantic type assignment. In addition, a set of inquiry and inquiry responses concerning speech acts are present. The generation process then proceeds as follows. The general multilingual network is entered and when decisions concerning MOOD are made, the grammatical features 'indicative', 'interrogative', and 'polarity' are made irrespective of which language is being generated. Since no language restrictions are given outside of the realization statements for the definitions of these systems, there is no difference internally in the representation of these systems to the state of affairs in the monolingual Penman generator. The inquiry specifications in the SPL (i.e., :command-q, :statement-q, :question-q, and :polarity-variable-q) are sufficient to motivate grammatical feature selection. Naturally, however, depending on which language is being generated, *distinct* sets of realization statements will have been collected. The pool of constraints following selection of these three grammatical features is as follows for each language:

English:	+Mood; +Subject; +Finite; Mood (Subject, Finite); Finite^Subject
Chinese:	none
Japanese:	+Particle; Particle = *ka;* Idc / Process; Idc^Particle

This simply follows the constraints as specified according to grammatical feature and language in Figure 3.

Subsequently, however, different states of affairs hold in the three languages. If the language is English, there are no further systems to consider; if, however, the language is Chinese or Japanese, then further options must be resolved. For example, if the :expectation-q inquiry (which we use to motivate the grammatical alternation under POLARITY TYPE: Section 3) calls for the grammatical feature 'particle' to be selected, the realizational constraints added are:

Chinese:	+Interrogator; Interrogator = *ma;* Interrogator ^ $
Japanese:	Particle ^ $

Combining these particular contributions of constraints with those that are generated by the other components of the grammar network for each of the three languages, i.e., those concerned with tense, aspect, process types, participants, lexical selection, etc., we obtain the very different clause structures shown in Figure 7. Thus, we can see how multiple generation possibilities which differ widely in structural properties can be generated from a single semantic specification and with a single multilingual network.

ister, is also being addressed using this extended framework, although this goes beyond the scope of the present paper.

[4] For more details of the currently implemented algorithms, see [Bateman *et al.*, 1991]

[5] This is partial in that we have omitted much information that is necessary for generation in any language but which is not central to the discussion here.

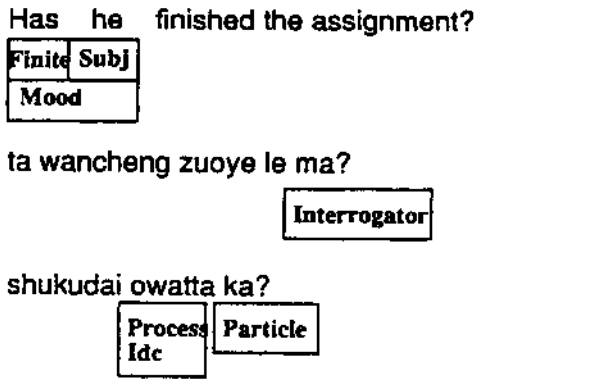

Figure 7: Distinct resultant structures for the input specification

5 Conclusions

Our trial languages, English, Chinese and Japanese are typologically very different structurally. For instance, Japanese has fairly elaborate verb structure, whereas Chinese does not; Japanese is classified as an SOV language, but English and Chinese are basically SVO, while German is SOV only in 'dependent' clauses. However, we have now illustrated how such languages can share systemic potential as well as diverge within a system network. We used genetically unrelated languages in order to show that the crucial criterion for grasping commonality is *functionality.* It can be noted that already for the restricted area of MOOD grammar, the functional viewpoint allows groupings of languages along dimensions of similarity and dissimilarity where a structural typology would not have suggested a basis for comparison. It should also be clear that many further languages could be added without substantial alteration: the important observations to make are that resources of a more general nature are shared and diversification in function across languages only appears later, at more delicate levels in the network, and that functional organization differs less than syntagmatic organization.

References

[Bateman *et ai,* 1991] John A. Bateman, Christian M.I.M. Matthiessen, Keizo Nanri, and Licheng Zeng. Multilingual text generation: an architecture based on functional typology. In *International Conference on Current Issues in Computational Linguistics,* Penang, Malaysia, 1991. Also available as technical report of the department of Linguistics, University of Sydney.

[Halliday, 1984] Michael A.K. Halliday. Language as code and language as behaviour: a systemic-functional interpretation of the nature and ontogenesis of dialogue. In Robin P. Fawcett, M.A.K. Halliday, Sidney Lamb, and A. Makkai, editors, *The Semiotics of Culture and Language,* volume 1. Frances Pinter, London, 1984. (Open Linguistics Series).

[Halliday, 1985] Michael A.K. Halliday. *An Introduction to Functional Grammar.* Edward Arnold, London, 1985.

[Kasper, 1989] Robert T. K asper. A flexible interface for linking applications to PENMAN's sentence generator. In *Proceedings of the DARPA Workshop on Speech and Natural Language,* 1989. Available from USC/Information Sciences Institute, Marina del Rey, CA.

[Li and Thompson, 1981] Charles N. Li and Sandra Thompson. *A reference grammar of Mandarin Chinese.* University of California, California, 1981.

[Mann, 1983] William C. Mann. The anatomy of a systemic choice. *Discourse Processes,* 1983. Also available as USC/Information Sciences Institute, Research Report ISI/RR-82-104, 1982.

[Matthiessen and Bateman, 1991] Christian M.I.M Matthiessen and John A. Bateman. *Text generation and systemic-functional linguistics: experiences from English and Japanese.* Frances Pinter, London, 1991.

[Matthiessen, 1985] Christian M.I.M. Matthiessen. The systemic framework in text generation: Nigel. In James D. Benson and William S. Greaves, editors, *Systemic Perspectives on Discourse, Volume 1.* Ablcx, Norwood, New Jersey, 1985.

[Matthiessen, 1990] Christian M.I.M. Matthiessen. Lexicogrammatical cartography: English systems. Technical report, University of Sydney, Linguistics Department, 1990. Ongoing expanding draft.