Multilingual Generation for Museum Applications

Alessandra Novello and Charles Callaway

ITC-irst, via Sommarive 18, Povo (Trento) Italy I-38050
{novello,callaway}@irst.itc.it
WWW home page: http://tcc.itc.it/

Abstract. Natural Language Generation (NLG) systems in English have been well established for the last two decades. Capabilities have been developed for creating scientific explanations of objects, automated technical documentation, intelligent tutoring systems, story generation, and a number of other application areas. Multilingual generation has been tackled in the last decade, with new systems arriving that produce texts in multiple languages simultaneously from a single knowledge representation. We describe the infrastructure for such a system that produces text in both English and Italian, where the Italian grammar and sentence planner were produced in-house at ITC-irst by modifying existing English components. We also briefly describe the overall structure of a multilingual generation system as well as prototype systems we have developed that use these new resources in a variety of ways.

1 Introduction

Multilingual natural language generation (NLG) from a single source representation is one of the holy grails of NLG research. Current NLG systems have complex, hand-coded grammars and other components which require substantial time and expertise to create and maintain. Thus an NLG system consisting of a set of parallel generators, each working separately on a single language, represents a waste of costly resources. Given that the knowledge bases that generation systems use as input are *interlingual* (or language-neutral) in nature, it makes sense to include only the necessary language-specific architectural modules while utilizing as many language-independent components as possible. Thus one important area of study in NLG is finding out what tasks are inherently language specific.

As in parsing, the English language has long dominated NLG research. But recent systems, described below, have taken the highly language-dependent tasks of sentence planning (deciding on sentence-level mappings between predicate arguments and their sentential roles) and surface realization (involving grammars, linear precedence, and morphology) and produced systems capable of generating a number of Western and Eastern European languages. However, there is still no general-purpose sentence planner or surface realizer for Italian.

We thus describe a new Italian sentence planner and surface realizer developed at ITC-irst which are based respectively on the FARE sentence planner [Cal 1995] and the popular FUF/SURGE generation system for English [Elh 1991]. The surface realizer is implemented as a systemic functional grammar whose inputs are hybrid semantic/syntactic descriptions of individual sentences. Furthermore, we briefly describe the STORYBOOK system [Cal 2002], an NLG research system for generating multipage, multilingual text. We have adapted STORYBOOK to use the new Italian sentence planner and surface realizer along with their English counterparts to create a multilingual NLG system.

Finally, we briefly describe several prototype applications using these new NLG resources in a museum context. While unfinished, these prototype projects illustrate potential benefits and problems of implemented NLG systems. We discuss what effects these may have on future systems and provide suggestions for future research directions.

2 Multilingual Natural Language Generation

Multilingual generation is becoming an increasingly important aspect of implemented systems that showcase the abilities of generation systems, especially in domains where readers can potentially be of any nationality. Most multilingual NLG systems require multiple sentence planners and surface realizers, one for each target language which must be deployed. Multilingual generators [Bat 1995,Par 1995,Ste 1996,Sco 1999,Cal 2003a] hope to leverage reusable components to produce texts in multiple languages with substantially less work than implementing an equivalent number of monolingual template or deep generators.

However, to be used effectively and efficiently in a multilingual generation system, these reusable components must be designed from the start for that purpose. A large array of resources are needed, which can function regardless of the language selected, such as revision and sentence planning rules, lexica, and pronominalization strategies.

These resources are needed because NLG systems generally consist of a sequence of architectural modules called a *pipeline* (Figure 1), where each module handles a different set of functions. For instance, discourse planning takes domain knowledge structures, selects which elements are most important, and organizes them into a specific discourse sequence which will convey the intended message. Sentence planners take individual elements of these semantic and pragmatic structures and assigns particular sentential roles to them such as subjects, objects, relative clauses, etc. Finally, the surface realization module is responsible for converting them into the actual surface text that the museum visitor reads or hears. Each of these elements must be structured to handle multilingual texts if the overall result is to be coherent to speakers of different languages. Additional modules may prove necessary to handle particular linguistic phenomena. For instance, revision and pronominalization modules refine and improve the quality of paragraphs composed of many sentences.

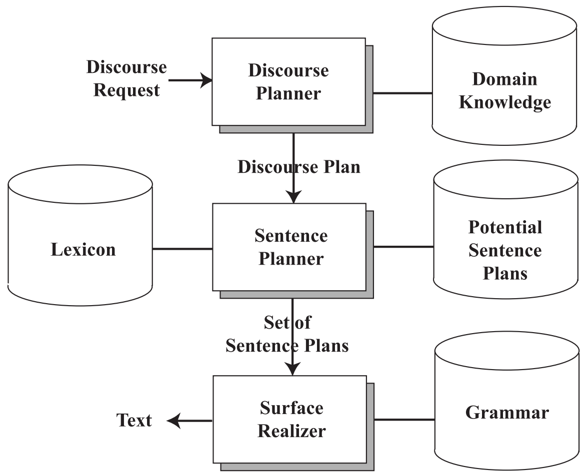

Fig. 1. A Typical Pipelined NLG Architecture

Individual applications which want to use natural language in their interactions with users must then create the types of domain structures that can be handled by the high-level discourse planner. Typically these are subgraphs of knowledge bases in the knowledge representation language of well-known KR formalisms such as KL-ONE [Bra 1985]. By passing a set of graph structures representing the desired communicative goal(s) in such a representation language, the discourse planner and subsequent NLG modules can organize them into a suitable text to be presented to the user. Thus, for example, if the user wishes to hear information about a painting, the name of the painting, its author, date of creation, *etc.* can be sent to the discourse planner in propositional form along with additional parameters such as communicative intent, the user's language, and size constraints to produce a text describing that painting. The resulting text, whether in HTML, ASCII, or whatever form, can then be displayed to the user in a browser, word processor, etc. or sent to a text-to-speech synthesizer.

3 Multilingual Sentence Planning and Lexicon

While a discourse planner should be inherently multilingual (*i.e.*, the internal structure of documents typically remains unchanged even after translation), a sentence planner is most certainly not. While a plan for an individual sentence does contain some language-independent concepts such as semantic roles (agent, patient, etc.), other elements are highly variable according to language: passivization rules, clefting constraints, verb-governed adjunctivization, etc.

Other modules between the discourse planning and sentence planning levels may or may not be language-dependent. For example, pronominalization rules for English and Italian are different, while newness/givenness features don't seem

```
(FINDING-IN-LOCATION
  (ENGLISH (CASE-FRAME (PROCESS-RESULT RESULT-REL RESULT RESULTING NO-ROLE))
           (LEX-VERB (((LEX "find") (TYPE COMPOSITE) (REFLEXIVE YES)
                       (RELATION-TYPE LOCATIVE))))
           (LEX-VERB-PREPOSITION ("in")))
  (ITALIAN (CASE-FRAME (PROCESS-RESULT RESULT-REL RESULT RESULTING NO-ROLE))
           (LEX-VERB (((LEX "trovare") (TYPE COMPOSITE) (REFLEXIVE YES)
                       (RELATION-TYPE LOCATIVE))))
           (LEX-VERB-PREPOSITION ("in"))))

(USER
  (ENGLISH (LEX-NOUN (((CAT PERSONAL-PRONOUN) (PERSON SECOND)
                       (NUMBER SINGULAR)))))
  (ITALIAN (LEX-NOUN (((CAT PERSONAL-PRONOUN) (PERSON SECOND)
                       (NUMBER SINGULAR))))))

(STATUE-ROOM
  (ENGLISH (LEX-NOUN (((CAT PROPER) (LEX "Statue Room")))))
  (ITALIAN (LEX-NOUN (((CAT PROPER) (LEX "sala delle statue"))))))
```

Fig. 2. Lexicalization of various concepts in both English and Italian

to be. Languages can be to either a greater or lesser extent similar as well: both Chinese and Italian allow zero anaphora, for example.

To create a usable multilingual generation system that would allow us to produce either English or Italian texts in our various research projects, we began the process of porting to Italian the language dependent components of our generation system StoryBook [Cal 2002]. StoryBook contains components for pronominalization, sentence planning, revision and surface realization, among others, but at present only sentence planning and surface realization have been attempted.

The sentence planner contains two parts: a method for mapping elements of the discourse plan (consisting of ordered sets of knowledge-based concepts) into sentence-sized groups with distinct semantic roles, and the lexicon, which contains language-specific information that relates individual concepts to raw linguistic data. Because this linguistic data is tightly linked to the surface realizer, the lexicon holds information in the representation formalism it expects. Thus the sentence planner can be seen as a function that maps knowledge-based concepts into linguistic-based tree structures, where pieces of each tree are constructed from the cumulative results of lexicon accesses.

To illustrate this point, the concept `FINDING-IN-LOCATION`, which is lexicalized by the verb "to find (in)" in English and by the reflexive verb "trovar-si (in)" in Italian, is depicted in Figure 2. The linguistic representation is that of *functional unification*, described in the following section. In our example, both lexical items have the semantic process type locative and the second argument under the category Prepositional Phrase; the Italian verb, further, has the fea-

```
((cat clause)
 (proc ((type composite)
        (relation-type locative)
        (reflexive yes)
        (lex "trovare")))
 (partic ((affected ((cat personal-pronoun)
                     (person second)
                     (number singular)))
          (located {^ affected})
          (location ((cat pp)
                     (prep ((lex "in")))
                     (np ((cat proper)
                          (lex "sala ...
```

```
((cat clause)
 (proc ((type composite)
        (relation-type locative)
        (reflexive yes)
        (lex "find")))
 (partic ((affected ((cat personal-pronoun)
                     (person second)
                     (number singular)))
          (located {^ affected})
          (location ((cat pp)
                     (prep ((lex "in")))
                     (np ((cat proper)
                          (lex "Statue ...
```

Fig. 3. Italian (left) and English examples composed from elements in Figure 2.

ture (reflexive yes), which entails, at the grammar level, the presence of the clitic reflexive particle "-si" and its correct treatment dependent on the verbal mood.

In addition, Figure 2 contains examples of two other concepts representing noun structures. Thus for example with the proposition (USER FINDING-IN-LOCATION STATUE-ROOM), the sentence planner can retrieve appropriate lexicon items according to language and construct a sentence-level linguistic representation from smaller linguistic representations, along with adding additional features such as definiteness which can be computed from the discourse context. The surface text of this proposition could then be realized in English as "You find yourself in the Statue Room," or in Italian as "Ti trovi nella sala delle statue." Once this stage in the generation process has been concluded, the linguistic representation is passed to the surface realizer, described next.

4 The FUF/SURGE Generation Grammar for English

While the multilingual descriptions for lexical items are still partly language-dependent, the second major component, the FUF/SURGE surface realizer package, is completely language dependent, and thus our work on grammar and on morphology required much more time than that for sentence planning.

SURGE (Systemic Unification Realization Grammar of English) is one of the most comprehensive computational grammars of English available for NLG. It is implemented in FUF (Functional Unification Formalism), a special programming language which operates over *functional descriptions*, which are recursive matrices of attribute-value pairs containing a mixture of semantic and syntactic specifications (Figure 3). SURGE is a functional unification implementation with a heterogeneous approach: it mostly includes rules of systemic grammar, but also of lexicalist grammars such as HPSG and the Text-Meaning Theory.

Upon receiving a functional description (FD) representing a sentence to be generated, the FUF unification system compares the FD to the SURGE grammar to ensure its grammaticality. Simultaneously, default constraints and linear precedence information are added, which enables the surface realizer to add closed-class lexical items and determine linearized constituent order. Finally, morphology and punctuation rules are applied to ensure the resulting sentence corresponds to the norms of written text.

FUF/SURGE is a *generalized* surface realizer, in the sense that it does not force users to use a particular knowledge source for generating its input functional descriptions (the PENMAN system [Bat 1995], for example, requires the use of an accompanying knowledge representation language). Thus the system can be easily combined with other existing components to create a complete, customized NLG system. Because there was no generalized surface realizer for Italian, we devoted resources towards creating and testing one in a series of prototype museum applications.

5 Porting FUF/SURGE to Italian

To enable the creation of applications that offer multilingual applications of text generation, we began the process of porting the SURGE grammar from English to Italian [Nov 2003]. Because the surface realizer is only one architectural component of a much larger NLG system, it was important that the Italian surface realizer match its English counterpart. Thus throughout our efforts, we strove to make the input representation as similar as possible to the existing SURGE test suite.

The porting process itself involved changes to morphology, the linearizer and the grammar. We first worked on correctly generating individual Italian words by adding inflection and features such as gender which are typically lacking in English; we next controlled the interaction between words and added prepositional and verbal contractions. Finally, we began work on the grammar, which required more time than the changes just described.

During the creation of the new Italian grammar, we kept the same high-level organization of SURGE, which is modularized as a set of sub-grammars, each of which is in turn divided into systems. We modified each system in two ways: adding constraints and features to grammar branches and changing the linearization patterns to reflect the differences in linear precedence between English and Italian. We also made changes at the lexical level by modifing syntactic and thematic structures of some lexemes.

As result of this preliminary Italian grammar, we obtained enough coverage (a measure of how many syntactic constructions can be correctly generated by the system vs. how many are actually present in the language) to allow for the production of multiple paragraphs of simple text. The end result of this process was an Italian grammar that, while not capable of generating sentences of exceedingly complex syntax, was capable of sufficient coverage to operate in an application in a museum domain. A description of this porting process has been quantitatively illustrated in [Nov 2003].

Further work we have conducted since then includes generation of Italian sentences with clitics, indicative versus imperative moods, adjuncts and disjuncts in non-finite moods, co-occurance of multiple adjuncts, and increased word-order flexibility.

6 Prototype Applications

NLG techniques have been used in various projects aimed at generating extended descriptions and multimedia presentations of museum exhibits. For example, the HIPS project [Not 1998] used mobile computers and localization hardware to produce contextualized audio descriptions, while the recent MPIRO project [And 2002] used multilingual, multiparagraph NLG and a browser-based interface to describe a museum collection in English, Italian and Greek, making contrasts and comparisons to objects that had been previously described.

However, NLG applications can provide other information besides extended explanations of museum objects. One class of museum applications can best be characterized as "post-visit summaries" where a visitor's route through the museum has been tracked with localization hardware. These summaries can take the form of automatically generated reports for a museum staff of aggregate museum visitor behavior or even personalized "postcards" describing a visitor's museum trip to a relative back home. Both of these applications would require logging capability as well as algorithms for the interpretation of that logged data. In each case, a discourse planner must be equipped to understand this interpretation, but could then generate the reports or postcards using standard NLG methods.

A second application when visitors have mobile computing devices is to allow the museum visitor to request directions to other museum exhibits of interest. Direction giving would require that the visitor have an interface mechanism to indicate where she wants to go, a planning mechanism to decide which route to take, and then a discourse/strategy module for deciding the sequence of directions to give as well as when they should be given and how many pieces of information should be given at each point in time. Because this last stage should return directions in either spoken or written text, natural language generation techniques are necessary.

A much more complex application for museums would be the general category of "service dialogues". This would include mainly full-scale dialogue systems which could allow visitors to interactively ask for more information about particular exhibits or even interact with a virtual tour guide. Whether they use spoken or written text, the underlying method of creating responses for the visitor's questions in such cases would be multilingual NLG technology.

7 Conclusions and Future Work

Using a standard surface realizer for English as a starting point, we have created a significant new resource: a general-purpose surface realizer for Italian which can be integrated into an existing NLG system. In addition, we have begun adding multilingual generation support to various prototype applications currently being developed at ITC-irst.

There are two major foci for future work. First, we must show that the new Italian surface realizer is capable of substantial enough coverage. Because demonstrating this by hand is a time-consuming process, we hope to turn to automatic methods instead, such as those based on TreeBanks [Cal 2003b], which can quickly determine coverage based on many thousands of example sentences. Second, we will need to test the generation system across a variety of domains and applications such as those described in the preceding section.

References

[And 2002] I. Androutsopoulos, J. Calder, E. Not, F. Pianesi and M. Roussou: MPIRO: Coherence, Relevance and Appropriateness in Information Presentation. In Proceedings of the 2002 CLASS Workshop. Verona, Italy, 2002.

[Bat 1995] John A. Bateman: KPML: The KOMET-Penman (Multilingual) Development Environment, Release 0.8. Technical Report. Institut für Integrierte Publikations- und Informationssysteme (IPSI). GMD, Darmstadt, 1995.

[Bra 1985] R. J. Brachman and J. G. Schmolze: An overview of the KL-ONE knowledge representation system. *Cognitive Science* 9(**2**):171–216, 1985.

[Cal 1995] Charles B. Callaway and James C. Lester: Robust Natural Language Generation from Large-Scale Knowledge Bases. In Proceedings of the Fourth Bar-Ilan Symposium on the Foundations of AI. Jerusalem, Israel, pp. 96–105, 1995.

[Cal 2002] Charles B. Callaway and James C. Lester: Narrative Prose Generation. *Artificial Intelligence* 139(**2**):213–252, 2002.

[Cal 2003a] Charles Callaway: Multilingual Revision. In Proceedings of the 9th European Workshop on NLG. Budapest, Hungary, April 2003.

[Cal 2003b] Charles B. Callaway: Evaluating Coverage for Large Symbolic NLG Grammars. In Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence. Acapulco, Mexico, August 2003.

[Elh 1991] Michael Elhadad: FUF: The Universal Unifier User Manual Version 5.0. Technical Report #CUCS-038-91, Dept. of Computer Science, Columbia University. 1991.

[Not 1998] Not, E., Sarini, M., Stock, O., Strapparava, C. and Zancanaro, M.: Information Adaptation for Physical Hypernavigation. In Proceedings of the I3 Annual Conference. Nyborg, 1998.

[Nov 2003] Alessandra Novello and Charles Callaway. Porting to an Italian Surface Realizer: A Case Study. In Proceedings of the 9th European Workshop on NLG. Budapest, Hungary, April.

[Par 1995] Cécile L. Paris, Keith Vander Linden, Markus Fischer, Anthony Hartley, Lyn Pemberton, Richard Power and Donia Scott: A Support Tool for Writing Multilingual Instructions. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence. Montréal, Canada, pp. 1398–1404, 1995.

[Sco 1999] Donia R. Scott: The Multilingual Generation Game: Authoring Fluent Texts in Unfamiliar Languages. In Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence. Stockholm, Sweden, 1999.

[Ste 1996] Manfred Stede: Lexical Semantics and Knowledge Representation in Multilingual Sentence Generation. Ph.D. Thesis, University of Toronto. Toronto, Ontario, 1996.