Jo Calder, Alexander C. Melengoglou
Division of Informatics, University of Edinburgh (UK)

Charles Callaway, Elena Not, Fabio Pianesi
TCC Division, ITC-irst, Trento (Italy)

Ion Androutsopoulos[*], Costantine D. Spyropoulos
Institute of Informatics and Telecommunications
NCSR "Demokritos", Athens (Greece)

Gerasimos Xydas, Georgios Kouroupetroglou
Department of Informatics and Telecommunications, University of Athens (Greece)

Maria Roussou
Foundation of the Hellenic World, Athens (Greece)

MULTILINGUAL PERSONALIZED INFORMATION OBJECTS

Abstract. The M-PIRO project targets the concept of personalized information objects ─ that is, entities capable of responding to requests for information by taking into account what the requester already knows, what they are most interested in, and how the related information is to be made available. M-PIRO's technology allows textual and spoken descriptions of exhibits to be generated automatically from an underlying language-neutral database, existing free-text descriptions, or a mixture of the two. The resulting descriptions, produced in three different languages (English, Greek, and Italian), are tailored according to the user's interests, background knowledge, and language skills. Particular research emphasis is placed on user modeling for personalization and authoring tools that allow museum caretakers to create on-line and virtual presentations of exhibits. Improvement of the synthetic speech output quality has been achieved via a closer integration between text generation and speech synthesis, and by using domain-specific trained prosodic models.

[*] Current affiliation: Dept. of Informatics, Athens University of Economics and Business, Greece

1. INTRODUCTION

In a world of burgeoning information resources we need evermore responsive and intelligent ways of locating and delivering them. The M-PIRO project (Multilingual Personalized Information Objects) aims at bringing together leading research work on natural language generation (NLG) and speech synthesis, with multimedia database, information management and virtual reality technologies developed for public access to museum and gallery resources. M-PIRO targets the concept of personalized information objects ─ that is, entities capable of responding to requests for information by taking into account what the requester already knows, what they are most interested in, and how the related information is to be made available.

The chosen context is cultural heritage: the project works closely with museums, galleries, and other "memory institutions" to develop technology responsive to their special needs. Museum missions, in fact, are changing: from the traditional, curatorial concern for scholarly understanding of objects and their conservation, the focus is shifting towards the visitors' interpretation of the objects in their historical and cultural context and towards making them widely accessible. The M-PIRO project provides tools and techniques to help this interpretative process by facilitating information delivery in a number of different languages and in a personalized way. It allows textual and spoken descriptions of exhibits to be generated automatically from an underlying language-neutral database, and from existing free-text descriptions. The resulting descriptions, generated in three different languages (English, Greek, and Italian), are tailored according to the user's interests, background knowledge, and language skills. For example, simpler vocabularies and syntactic structures may be used with children than with adults and/or experts.

A prototype has been implemented capable of generating descriptions of 50 archaeological exhibits of different kinds, which shows the potential of the developed architecture for the delivery of information presentations at public access kiosks in museums, or through the web (Figure 1). The interaction scenario has also been extended to an immersive virtual reality environment, where another prototype delivers descriptions in spoken form (Figure 2).

2. DYNAMIC GENERATION OF MULTILINGUAL PERSONALIZED DESCRIPTIONS

One starting point of M-PIRO is the text generation facilities for English developed within the ILEX (Intelligent Label Explorer) project [Oberlander *et al*. 1998; O'Donnell *et al*. 2001]. ILEX was capable of generating web-based descriptions of diverse exhibits of modern jewellery located in the Royal Museum of Scotland. ILEX is particularly well-adapted at generating facts about collections of objects, such as museum catalogues. In addition to novel features such as clause aggregation, ILEX marks individual facts with importance and interest scores, allowing the text planner to create useful exhibit descriptions despite variations in the underlying data for different exhibits. Within M-PIRO, these facilities have been extended along a number of dimensions:

Figure 1: An example of the M-PIRO web interface customized for an English adult

Multilinguality. M-PIRO supports multiple languages: English, Greek and Italian. The production of linguistic resources for the latter two has been pursued with the aim of reusing work done on a systemic grammar for English [O'Donnell 1996] as much as possible. The ultimate goal was to experiment with an approach to multilingual generation that leverages grammar commonalities across different languages, allowing for faster resource development and easier maintenance [Bateman *et al*. 1991; Callaway *et al*. 1999]. Many languages share the same basic systems, when seen functionally, but differ in their realization, or in the way specific predicates or relations are phrased. A core set of grammatical systems has been identified starting from ILEX's English resources. Then, language-specific Italian and Greek [Dimitromanolaki *et al*. 2001a] resources have been defined as complementary grammatical subsystems, lexicons and micro-planning expressions.

Integrating existing text and automatically generated text. M-PIRO's interaction scenario and domain (the presentation of information to visitors of virtual archaeological museums) require a flexible integration of deep generation with shallow approaches. Indeed, much interesting information about relevant objects is already available to museum curators in the form of narrative text, anecdotes, or other complex texts that would be too difficult or expensive to generate from scratch. M-PIRO investigated this issue along various alternative directions. The implemented systems allow the flexible incorporation of existing pieces of text by associating (tri-lingual) strings to relevant domain entities and annotating them with information about their appropriateness with respect to user and discourse context. A template-based mechanism then allows the combination of strings with dynamically generated referring expressions or other strings, while the text planner guarantees that the resulting text is appropriately placed.

In parallel, a more substantial integration of NLG and Adaptive Hypermedia (AH) technologies has been investigated to make the best of the strengths of each — *e.g.*, by resorting to AH functionalities to convey information that could not be easily deployed by the deep generation component [Not and Zancanaro 2001]. The adopted AH technology is based on the Macronode approach [Not and Zancanaro 2000], which allows for the segmentation and annotation of pre-existing texts (as well as images and audio messages) into small content units. In this approach, the units can be reused in different discourse contexts after adjusting them dynamically to merge with the rest of the presentation.

Personalizing text with user modelling. Advanced user modelling facilities have been developed to support the personalization task by keeping track of what visitors see and the information that has already been presented to them. By refining previous models of users' interests and knowledge, we can produce descriptions that better match their needs. We can also arrange our descriptions so that they draw connections between previously seen exhibits, convey general information about styles or periods, and offer pointers to other related exhibits, thus helping the visitors obtain a more cohesive view of the presented information. M-PIRO's user models are more elaborate than ILEX's in that they contain information that allows the system to tailor not only the content but also the surface form of the resulting texts. It is thus possible to generate, for example, shorter sentences with simpler words when producing texts for children. We store the user modelling information in the personalization server (developed at "NCSR Demokritos"), a database that is interrogated by the system during generation and dynamically updated after each text is produced. The personalization server stores both static parameters associated with user types (e.g., the appropriateness of language expressions per user type) and dynamic information pertaining to individual users (e.g., the degree to which they have assimilated various pieces of information).

Authoring of domain and linguistic resources. To facilitate the specification of the domain-dependent knowledge and linguistic resources, which typically represents one of the major bottlenecks in developing and porting a generation system to different domains, an authoring tool (developed at NCSR "Demokritos") was created that enables museum curators to update the information source from which the texts are generated, and control the language and content of the resulting descriptions [Androutsopoulos *et al*. 2001]. Unlike systems such as KPML [Bateman 1997], the authoring tool is not intended to assist language technology experts in creating and maintaining domain-independent linguistic resources, such as large-scale grammars. In that sense, M-PIRO's authoring tool is closer to the symbolic authoring facilities of DRAFTER [Paris *et al*. 1996] and GIST [Power and Cavallotto 1996].

M-PIRO allows the author to modify not only the assertional part of the database ─ *i.e*., its entries ─ but also the schema of the database. This is a consequence of the fact that M-PIRO targets a broader range of domains; actually, any domain where descriptions of physical objects are needed. Besides museum collections, this extends, for example, to sales catalogues. Additionally, M-PIRO allows the authors to manipulate not only the information in the database, but also the domain-dependent resources that control how this information is rendered in natural language. This allows them to control, for example, the vocabulary and form of the generated sentences, as well as the structure of the resulting texts.

Integrating text generation with speech synthesis. Finally, M-PIRO has developed a state-of-the art multilingual Concept-to-Speech system. Novel techniques are being used to increase the naturalness of speech by tightly integrating speech synthesis with the NLG module. Synthesizers for English and Italian are constructed using Festival [Taylor *et al*. 1998], a common environment allowing the incorporation of newly developed resources; in our case for Italian [Cosi *et al*. 2001]. Greek speech synthesis is hosted by DEMOSTHeNES Speech Composer [Xydas and Kouroupetroglou 2001a], a research tool for speech synthesis.

M-PIRO aims at identifying the linguistic and discourse phenomena that significantly affect the intonational and prosodic contour of utterances, in order to allow the text generator to pass on this information to the speech synthesizers. The relevant information is delivered by the generation module through an XML-based annotation scheme. The quality of the audio output is of the utmost importance with immersive virtual reality environments, an option which M-PIRO actively explored in a prototype where users directly manipulated 3D objects to get contextualized and personalized audio descriptions (Figure 2).

3. TEXT GENERATION

Natural Language Generation (NLG) aims to develop computer systems that can effectively apply linguistic knowledge such as semantics, morphology and syntax, as well as knowledge of the application domain, to produce meaningful texts in human languages, based on non-linguistic data [Reiter and Dale 2000]. The generation process in NLG systems typically consists of the following four main stages:

- *Content Selection*, in which the system selects those pieces of information from its database that are most appropriate for a particular communicative purpose, user model etc,

Figure 2: The 3D collection of the M-PIRO virtual reality prototype: spoken descriptions are generated for each selected object.

- *Text Planning*, which specifies the structure of the text by ordering the facts to be conveyed and by establishing the relations that hold between them,
- *Microplanning*, which involves the processes of lexicalization (the choice of the particular words and word types, e.g. choosing the right verb, tense, voice, person etc.), aggregation (combining facts to achieve a more fluent and concise text), and referring expression generation (selecting appropriate phrases to refer to particular entities in the domain), and
- *Surface Realization*, which performs the mapping of the abstract specifications into natural text, including morphology, linear precedence, and formatting for display in web browsers.

3.1 The Exprimo text generation system

The generation system used in M-PIRO is Exprimo, a modular NLG system that adheres to the fairly standard pipeline architecture [Reiter 1994] of figure 3. In general, Exprimo contains a number of important advances over previous end-to-end generation systems. Among these are:

- The ability to generate in three languages using shared language resources
- The use of domain-specific linguistic input from an authoring tool
- Integration with user modelling
- The ability to change the text, both lexically and syntactically, depending on user type (personalization)

The domain model consists of the domain database and the domain semantics. The domain database provides specifications for the classes of the entities present in the domain (like ‘exhibit’, ‘god’, ‘historical-period’, and so on), as well as the entities themselves (such as ‘exhibit1’, ‘Apollo’, ‘Hellenistic-period’).

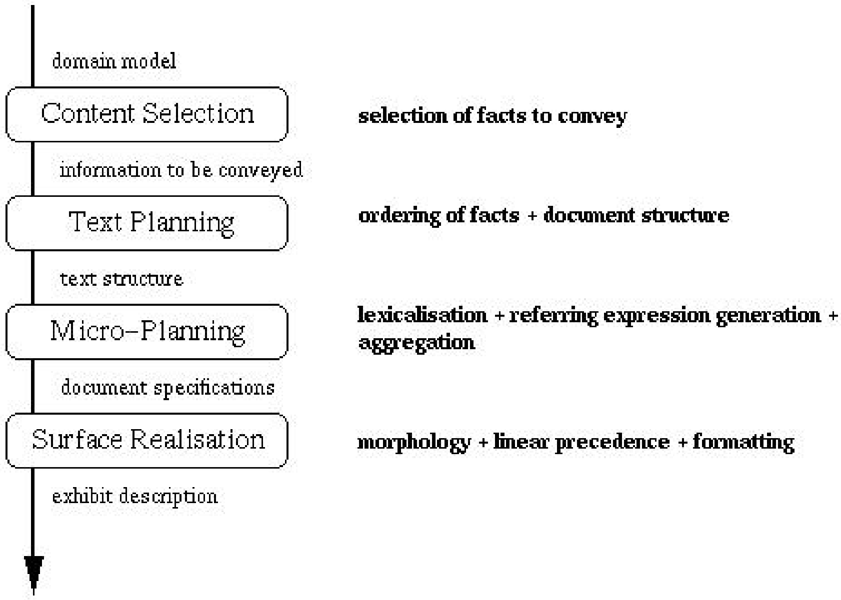

Figure 3: A generalized pipeline NLG architecture, to which Exprimo adheres

Each of the entities is described by means of a number of predicates, which are either attributes (*e.g.*, 'exhibit-height' takes values like '2.12.m') or relations that can exist between the entity and members of other classes in the domain (*e.g.*, the 'original-location' of an exhibit will be defined by another entity in the domain, such as 'Pergamos'). Typically, each class defines the set of predicates that are relevant to its members.

The domain semantics allows the system to interpret the domain database by providing predicate definitions and a class taxonomy. The predicate definitions specify domain restrictions and the way in which the predicates should be expressed. For example, the definition of the predicate 'current-location' will specify that this predicate can only be used to relate an instance of the basic class 'exhibit' to an instance of the class 'museum'. The taxonomy defines the hierarchy of classes in the domain. The *Content Potential* is an internal representation of the domain model in the form of a graph, reflecting the relations that exist between the entities in the database. Each instance of a relation or attribute in the domain (a fact) is a node of the graph (a fact-node), linked to the target domain objects (entity-nodes) by means of two arcs corresponding to the arguments of the predicate.

The Content Potential represents the information that the system can communicate to the user, and the way in which the individual pieces of information are connected to each other. When the system receives a user request for an entity description, the system makes the particular entity the 'global focus' of the text to be generated [Mellish *et al*. 1998]. The facts that are connected to the entity in global focus (i.e. those satisfying the criterion of structural relevance) become candidates for selection by the system. The facts that will eventually be selected by the system are most likely to be those best matching the criterion of inherent relevance as well, which is evaluated on the grounds of the assumed interest of the content to the user, the importance of the fact in question relative to the pedagogical goal of the system, and the user's level of assimilation of the information [Not *et al*. 2001.]

The Text Planning module builds the structure of the text by ordering the facts and defining the dependencies that exist between them. The task of the text planner is to build the text-structure, which is a tree whose leaves are the selected facts (in other words, each leaf is a mapping of a fact onto a text-node). The internal nodes of the text structure specify the exact order in which the facts should appear in the output text, and the discourse relations that hold between those facts. The microplanner takes this "rough" document plan and produces a detailed specification for the text, which is used by the realization module to produce the surface document. Microplanning includes the language-dependent task of lexicalization that involves not only the choice of the words to use, but also of the syntactic structures that will be used to express the content of the propositions.

Exprimo performs surface realization using its computational grammar for the specific target language: a systemic grammar [Halliday 1985] that consists of choice systems, each offering a set of alternatives representing grammatical features. Realization is carried out by means of a series of choices (traversal of the systemic network), where each choice is made at a different level in the hierarchy of ranks, starting from the clause-complex rank and ending at the word-rank, which specifies the exact grammatical and syntactic features at the lexical level.

3.2 Integration with the Authoring Tool

Exprimo is tightly integrated with a domain authoring tool. There are two main aspects to this: first, the domain descriptions produced using the tool can be stored as XML files which Exprimo can then use to generate text. Second, the authoring tool can use Exprimo to preview texts, giving immediate feedback to authors on the accuracy and appropriateness of new additions. The authoring tool supports user modelling, allowing information on the user types and specifications such as the desired number of facts (and thus page length) for each user type to be set. It is also used to set the interest, importance, and repetitions for each user type. The repetitions parameter defines how many times the information must be conveyed before we can consider that the user has assimilated it.

In general, the domain author can specify the interest, importance and repetitions of a field at the entity type (class) that introduces it. In the current domain, for example, the interest of the 'creation-period' field for adults is set to 3 at the 'exhibit' type, ensuring that all 'exhibit' entities inherit that value. The inherited interest, importance, and repetition scores can then be overridden, at individual entities.

The authoring tool also allows the author to specify microplans. These are expressions controlling the lexicalization of each field (e.g., which verb to use, in what tense and voice, etc.) and encoding preferences that affect aggregation and the generation of referring expressions (e.g., whether or not the resulting clause may be aggregated, the desired case of the referring expressions, etc.). Multiple microplans can be specified per field, and each microplan is assigned a numerical score that indicates its appropriateness for each user type. Preference is given to microplans with the highest appropriateness scores that have been used fewer times in documents generated for the particular visitor. This allows the author to control the language expressions per user type (e.g., use simple verbs in active voice when generating for children), while leading at the same time to greater variation in the surface form.

3.3 Aggregation

The Content Selection module packages information as “verb-based, clause-sized propositions, each of which is realised as a single sentence” [Shaw and McKeown 1997]. Texts composed exclusively of such single-fact sentences are very likely to contain repetitions and redundancies, and are almost certain to be considered boring and unnatural by human readers. To overcome this problem, NLG systems make use of aggregation algorithms, which combine semantically related propositions in order to produce a more concise and more coherent text.

Exprimo is designed to generate personalised descriptions, based on its model of potential user types. In this context, it is desirable that the aggregation module also possess the flexibility to adjust the syntactic complexity of the output text to the requirements of the target user type. The authoring facilities of M-PIRO allow domain authors to specify the maximum number of facts that the system should convey to a particular type of user in each sentence.

At the heart of Exprimo's aggregation module is a sophisticated algorithm that allows the system to examine the complexity and the quality of the constructed sentences in parallel. During each iteration of this algorithm, the system examines the effect of adding one more fact to the sentence being constructed on the quality of the sentence (provided that the sentence has not already reached the maximum size allowed by the user modelling parameter), as well as on the quality of the remaining (that is, of the still unaggregated) text. Thus, the system never sacrifices text quality in order to increase the length of a sentence.

The module is equipped with a set of standard aggregation operations (such as simple conjunction and syntactic embedding), each of which is able to express a certain type of rhetorical relation holding between two neighbouring facts in the text plan. Even though the module is capable of producing a significant variety of syntactic forms based on this set, as the length of a sentence increases, the module is more likely to employ one of its non-standard aggregation operations to resolve conflicts between sequences of incompatible standard operations. This may result in complex sentences using semicolons, qualifiers, etc. [Melengoglou 2002]

The following example illustrates the effect of different values of the user modelling parameter on a series of four facts generated by the system:

- *maximum facts per sentence = 1 (no aggregation)*
  This exhibit is a rhyton. It was painted by the painter of Sotades. It was painted with the red figure technique. Today it is located in the Musee du Petit Palais.
- *maximum facts per sentence = 2*
  This exhibit is a rhyton, painted by the painter of Sotades. It was painted with the red figure technique and today it is located in the Musee du Petit Palais.

or

  This exhibit is a rhyton. It was painted by the painter of Sotades with the red figure technique. Today it is located in the Musee du Petit Palais.
- *maximum facts per sentence = 3*
  This exhibit is a rhyton, painted by the painter of Sotades with the red figure technique. Today it is located in the Musee du Petit Palais.
- *maximum facts per sentence = 4*
  This exhibit is a rhyton; it was painted by a painter called the painter of Sotades with the red figure technique and today it is located in the Musee du Petit Palais.

3.4 Comparisons

As already discussed, the system stores the information that it can convey to the users about the entities in its database in the form of fillers (which are either values or pointers to other entities) of a number of predicates that are common to the members of a certain class. For instance, members of the class *exhibit* can have predicates such as *creation-period* and *original-location*, while members of the class *person* can have different ones like *person-country*, etc. Thus, each entity can be described by means of a finite set of predicates, and it is possible to group the members of a class in a number of ways, based on the fillers of their common predicates. For example, in the domain of museum exhibits, we may classify exhibits according to their origin (e.g. several exhibits originate from Attica, others from Argos etc.), according to the historical period during which they were created (Archaic, Classical, Hellenistic etc.), and so on.

Exprimo generates comparisons between the current and previously examined entities, thus helping the users of the system build a more complete picture of the concepts that are presented, make associations and reach valuable conclusions. By being the result of the user's own navigation, such comparisons result in texts that are more personalised and, therefore, more interesting, attractive and intelligent.

Two types of comparisons are performed, based on entity similarities and differences (in terms of the resulting discourse relations: RST-SIMILARITY and RST-CONTRAST). 'Contrast' comparisons always take place between the entity in current focus and a group of previously visited entities (never an individual entity).

This group is always directly related to the current entity; it belongs either to the same class as the entity, or to one of the entity's super-classes.

Central to the design of this module is the aim to avoid making individual, full-clause references to previously seen entities, such as: *"Like the amphora that was decorated by the painter of Kleofrades, this lekythos was created during the archaic period"*. Such past references both distract the users from their focus of attention (the current exhibit) and tend to lead to confusion, as users are unlikely to remember which of the (possibly several) exhibits that they have already seen is the one that actually matches a description. Thus, such comparisons tend to be boring, if not irritating, while their educational value is questionable. The comparison module used in Exprimo attempts to overcome this problem by using the type hierarchy to group previously examined exhibits into broader categories (*classes*) and make either group references (references to the class, if a comparison holds for all the previously mentioned members of the particular class), or short (name-only) individual references, when the system knows that the exhibit's name is sufficient to make a unique reference. In order to ensure that the generated comparisons are meaningful, the module performs comparisons only between entities belonging to the same sub-type of the class 'exhibit' (i.e. between vessels, statues, coins, and so on). For instance, the system will attempt to compare a panathenaic amphora with a hydria or a black kantharos (as they are all vessels), but never, say, with a drachma or with a suit of armour.

The following descriptions of seven vessels illustrate the use of comparisons in Exprimo:

> This exhibit is a white lekythos; it was created during the classical period and it dates from circa 440 B.C. […] This white lekythos was decorated by the Painter of Achilles. Scenes related to war were quite frequent in Athenian art during this period.
>
> This exhibit is a panathenaic amphora. *Like the white lekythos, it was created during the classical period*. […] This panathenaic amphora dates from 490 B.C. and it was painted with the black figure technique. It originates from Attica.
>
> This exhibit is a marriage cauldron; it was created during the classical period and it dates from between 420 and 410 B.C. […] This marriage cauldron was decorated with the red figure technique. *Like the previous vessels, this marriage cauldron was originally from Attica*.
>
> This exhibit is a hadra ware hydria, created during the hellenistic period, which ranged between 323 and 31 B.C. […] *Unlike the previous vessels, which were originally from Attica, this hadra ware hydria was originally from Rhodes*. Today it is located in the Archaeological Museum of Rhodes.
>
> This exhibit is a panathenaic amphora; it was created during the hellenistic period and it dates from 320 B.C. *Like the previous panathenaic amphora, it*

> *was painted with the black figure technique*. It is now exhibited in The John Paul Getty museum. […]
>
> This exhibit is *another* amphora, created during the archaic period. It dates from the early 5th century B.C. […] This amphora was painted by the painter of Kleofrades. *Unlike the previous amphoras, which were painted with the black figure technique, this amphora was painted with the red figure technique*.
>
> This exhibit is an aryballos; it was created during the archaic period and it belongs to the corinthian type. It is spherical in shape. *Like the panathenaic amphoras, this aryballos was decorated with the black figure technique*. […]

3.5 Forward pointers

An important aspect of the web-based presentation of objects is the generation of forward pointers, which suggest a number of alternative paths that the users can follow to examine other objects that are in some way related to the one in the current description.

Exprimo generates lists of forward pointers in all three languages dynamically, with the use of the generation engine. The Forward Pointers module locates those pieces of information (attributes of the current exhibit) that are assumed to be most relevant to the user (according to the user modelling assumptions made by the domain author), and builds lists of hypertext references to exhibits which have not yet been visited by the user, and which also possess the attribute in question.

The 'relevance' of each fact in this context is the value of the *inherent relevance* parameter (see section 3.1), and is thus not synonymous to *interest*. Apart from the user's assumed interest in the fact, this parameter also considers the *importance* of this information to the communicative objectives of the system, and the degree to which a fact is already known to the user (*assimilation* of the fact). This allows the Forward Pointers module to be more flexible whenever the user requests subsequent descriptions of the same object, as facts which are introduced to the user for the first time (i.e. unassimilated facts) have a higher probability of forming a list of forward pointers.

The following example presents a description of an ancient Greek coin, together with the generated lists of forward pointers, which are associated with particular facts in the description:

> This exhibit is another tetradrachm, created during the classical period. It dates from between 440 and 420 B.C. It has an image of Athena crowned with a branch of olive, her tree, on its obverse. On the reverse there is a picture of her owl. Like the previous tetradrachm, this tetradrachm is made of silver. It was originally from Attica.
>
> Other exhibits created during the classical period:

-A hydria painted by the painter of Meidias
-A lekythos decorated with the red figure technique

Other exhibits made of silver:
-A stater that originates from Croton
-A drachma originally from Attica

3.6 Generating Text with Adaptive Hypermedia Techniques: Macronodes

In many application domains, full natural language generation is not the most adequate solution due to inefficiency, high development costs, and the difficulty of formalizing knowledge ontologies. Often, pre-existing texts (and related images or audio commentaries) prepared by domain experts are a valuable knowledge source for assembling information presentations. Hybrid systems that combine both NLG and template techniques represent in these cases a promising solution [Reiter 1995].

Exprimo is a good candidate in this direction: it provides sophisticated NLG machinery to support generation from scratch (i.e., generation from a knowledge base representing entities and facts related by semantic relations) as well as mechanisms to support the integration of flexible templates. However, when large amounts of material must be reused, techniques from the field of Adaptive Hypermedia (AH) seem to be more effective in shaping and composing the existing information. Within M-PIRO we have investigated the benefits of complementing the Exprimo NLG system with a flexible approach to the composition of canned text such as the one supported by the MacroNode's adaptive hypermedia architecture, thus achieving a trade-off between information reuse and flexibility. In this enlarged scenario, Exprimo's NLG machinery allows for the generation of the most dynamic parts of the presentation, whereas the AH machinery allows for the enrichment of content and phrasing, while still allowing for personalization.

As a starting point for AH, we have considered the MacroNode formalism [Not and Zancanaro 2000] that has been designed with the purpose of building a new kind of flexible hypermedia system which can overcome the limitations of the current adaptive hypermedia and the difficulties/costs of dynamic hypermedia. In this framework, available multimedia material is fragmented into content units which are annotated with essential semantic and linguistic data (i.e. a description of the main topic of the content unit), the semantic relations existing with other units (in terms of RST rhetorical relations [Mann and Thompson 1987]), and the different ways in which the unit can be presented to the user.

A content unit is called a *macronode* and, for textual information, it typically corresponds to a paragraph. More generally, a macronode can correspond to data in different media, for example text, audio, images, or a combination of these. The macronode approach exploits text planning techniques and linguistic manipulation rules derived from the field of NLG when composing from content units. At generation time, the annotated data is processed by a Presentation Composer that uses some NLG text planning strategies to determine the most appropriate discourse organization for the macronodes relevant for the current discourse context. The module explicitly maintains a representation of the discourse history and is able to exploit a model of the user's interests and knowledge to best tailor the output.

The schema-based text planning phase, which relies on the rhetorical organization of the macronodes in the repository, guarantees that the composed message displays (i) a coherent rhetorical structure, (ii) a topic flow compatible with discourse and interaction context, and (iii) cross-references to material already presented as well as to the user interests. For example, discourse strategies are encoded to avoid presenting already known information, to choose the kind of information for which the user's interest is higher, and to present new information when the user goes back to a previous topic. In order to assure cohesion, each macronode also encodes different realizations of the same content. These alternative realizations are stored as a graph. Each node of the graph represents a minimal unit of content, whereas arcs contain the conditions under which the target node can be realized (given that the source node has been realized). Each path in the graph is a possible realization of the content of the macronode.

Macronodes can be seen as an implementation of portions of canned text whose composition and linguistic form can be adjusted according to the discourse and interaction context. The advantage of splitting a text into small content units is that it allows reusing the same material in different discourse contexts, and permits a more flexible integration with the rest of the presentation. By integrating the MacroNode approach into Exprimo, the final system would be able to (i) reuse existing texts that could otherwise not be easily generated by Exprimo, (ii) still guarantee that they are integrated coherently in the overall presentation and partially adapted to the user needs, and (iii) opportunistically select the most appropriate generation machinery according to the type of information to be conveyed. Although, from the implementation point of view, the macronode approach has not been completely integrated into Exprimo, substantial headway has been made using two separate methods: constructing a new discourse planner that controls both Exprimo and the Macronode system as surface realizers, and adding discourse planning rules to Exprimo that would allow it to use macronodes as an alternate surface realizer [Not and Zancanaro 2001].

4 AUTHORING TOOLS FOR DOMAIN AND LINGUISTIC KNOWLEDGE

One of M-PIRO’s most ambitious goals has been to produce technology that is not only portable to new domains where object descriptions are needed, but ported to new domains solely by domain experts. By “domain experts” we mean people with expertise on the objects to be described, but with no background in language technology. Curators from the Foundation of the Hellenic World in Athens were involved from the beginning of the project in the formative evaluation, which helped define usability specifications for the authoring software.

The authoring tool is the main instrument for attaining portability. It allows domain experts, such as museum curators, to configure the system for a new collection of exhibits, or to modify the system’s knowledge of an existing collection by allowing them to manipulate the system’s domain-dependent information.

Resources of M-PIRO which are most important for the authoring tool include:

- *The Database Schema*: A structural description of the database, specifying:
  - The types of entities that are modelled by the database, such as artist, coin, statue, and vessel.
  - The sub-types and super-types of each entity type, in the form of a hierarchy. For example, coin, statue, and vessel are all sub-types of the entity type exhibit; vessel has sub-types like amphora and hydria, and artist has sub-types like sculptor and painter.
  - The types of relationships that entities of the various types may participate in. In the current domain, for example, there is a sculpted-by relationship between statue and sculptor entities, in that entities of the former type were sculpted by entities of the latter.
  - The attributes that entities of each type possess, such as the name and height of a statue entity. The current implementation of the authoring tool provides full control over all the parts of the database schema.
- *The Database Entries*: These provide information about particular entities. For example, there may be a statue entity whose name is “Doryphorus”, whose height is 2.12m, and which was ‘sculpted-by’ a sculptor entity called “Polyklitus”. The current implementation allows domain experts to insert, and later modify, database entries. Future work will also consider ways to import entries from existing databases, especially databases complying with museum standards.
- *Domain-dependent Linguistic Resources*: Many of M-PIRO’s linguistic resources, such as the grammars of the supported languages and the lexicons that contain closed-class words, like articles and prepositions, are domain-independent. Some linguistic resources, however, are domain-dependent, and need to be configured when moving to a new application domain:
  - The domain-dependent lexicon. This provides information about nouns and verbs that are relevant to the domain, including semantic information, such as the types of entities that a noun can refer to. The authoring tool provides facilities to manipulate the domain-dependent lexicon, which hide from the domain experts parts of the lexicon that require linguistic expertise. For example, there are mechanisms that infer automatically morphological features and generate all the forms of inflected words from their base forms.
  - Micro-planning expressions. These specify how relationships and attributes can be conveyed in the supported languages; for example, which verb to use to express a relationship, how the participants of the relationship relate to the subject and object of the resulting phrase, etc. The current implementation of the authoring tool provides full control over micro-planning expressions.
- *Partly Pre-generated Texts*: These are half-ready textual descriptions of objects, that are used when the information to be conveyed is too complex or too object-specific to express via full natural language generation. M-PIRO

addresses two kinds of partly pre-generated texts: parts of templates, which are almost completely canned pieces of text, and the macronodes described above. Facilities for authoring macronodes have been developed and tested separately.

M-PIRO's authoring tool (Figure 4) is intended to be used at two stages: domain authoring and exhibit authoring. Domain authoring defines the database schema and the domain-specific linguistic resources, while exhibit authoring provides the database entries and the partly pre-generated texts. Although the authoring tool is designed to be used by domain experts with no language technology expertise, some general training on the use of the tool is still required at both stages. It is possible to train only a few domain experts as domain authors who will be responsible for the initial configuration of the system and "advanced" modifications, such as the addition of new types of exhibits, or the tailoring of the domain-dependent linguistic mechanisms. The task of entering and maintaining information about particular exhibits can then be assigned to a larger number of exhibit authors, who may have received briefer training on the use of the authoring tool.

M-PIRO's authoring tool also provides a preview mechanism, which helps the authors detect anomalies in the content or realization of the generated texts, directing them to related parts of the database and domain-specific language resources. This can be seen as an attempt to link symbolic authoring to the WYSIWYM approach [Power and Scott 1998], where authors interact with the system entirely via generated texts that reflect both the content of the database and the options that are available to update it.

4.1 Domain authoring

The current implementation of the authoring tool provides full control over all the parts of the database schema, namely the hierarchy of the available entity types, relationships, and attributes.

The available entity types and their hierarchy are defined by manipulating the tree structure in the left pane of the database tab. In Figure 5, the tree structure shows that both vessel and statue are sub-types of the entity type exhibit, while amphora, hydria and rhyton are some of the sub-types of vessel. Types like exhibit, historical-period and place in Figure 5, that have no super-types, are called basic entity types; they are shown as children of the pseudo-node basic entity types in the tree structure of the authoring tool. Although M-PIRO's generation engine supports multiple inheritance (*e.g.*, entity types with more than one super-types), the tool only allows single inheritance, to make it easier to use by domain experts.

The tree structure initially contains only the basic entity type node and the built-in data-types. Basic entity types have to be linked to types of the Upper Model [Bateman 1992], a built-in domain-independent hierarchy of entity types. Linking a basic type to a type of the Upper Model in effect allows the generation engine to treat the basic type as a sub-type of the Upper Model type.

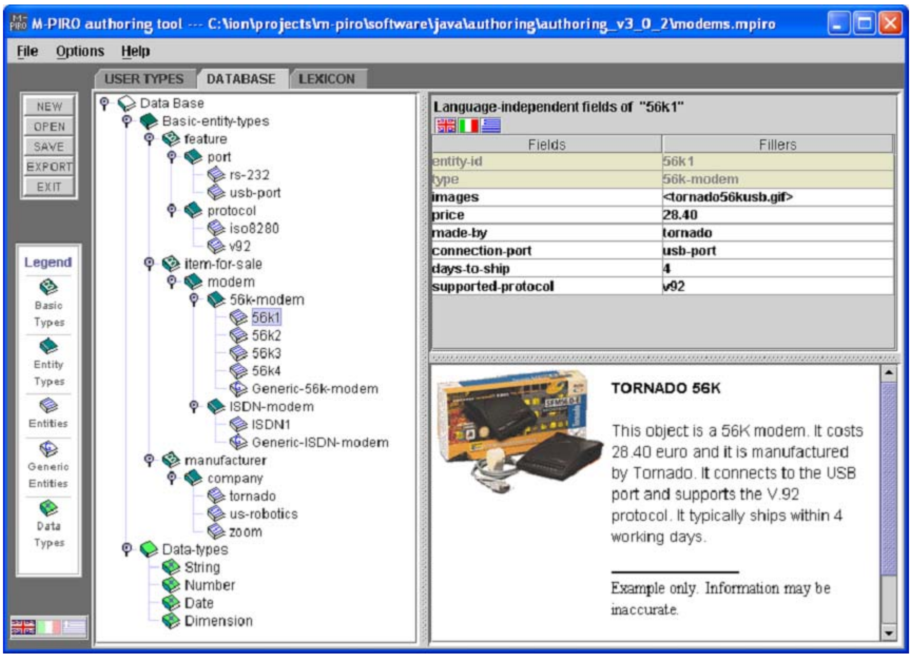

Figure 4: The M-PIRO authoring tool

A basic type can be linked to more than one types of the Upper Model; this is the only case where the authoring tool in effect allows multiple inheritance.

4.2 Domain-dependent linguistic resources

Each entity type can be associated with one or more nouns; for example, the entity type statue can be associated with the noun whose lexicon identifier is *statue-noun*, and whose English form is "statue". This allows the generation engine to use the nouns to refer to the corresponding entity types (e.g., "This is a statue."), and to generate appropriate noun phrases when referring to entities (e.g., "This statue was sculpted by… "). The associations between entity types and nouns are specified in the right pane of the database tab, as shown in Figure 5. The area below the table shows the nouns that are associated with the entity type that has been selected in the left pane, as well as nouns that are associated with super-types of the selected type.

New nouns can be added to the domain-dependent lexicon through the lexicon tab of the authoring tool. The tool encourages the domain authors to keep the lexicons of the supported languages aligned, by treating each lexicon entry as a triplet that contains information about equivalent nouns in the three languages.

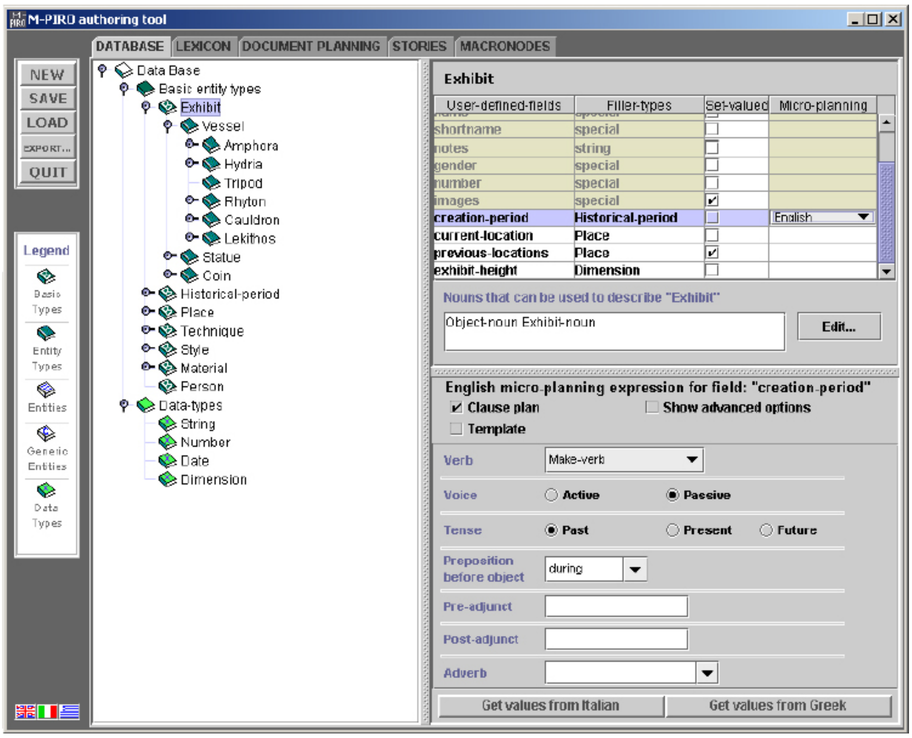

Figure 5: Domain hierarchy (left) and linguistic editing (right)

The internal form of the entries is more complicated: for example, the Greek part of the entry contains several features pertaining to the inflection of the noun, the position of the stress, etc. The authoring tool incorporates mechanisms that hide features of this kind that require linguistic expertise [Dimitromanolaki *et al*. 2001b].

M-PIRO's generation engine also incorporates morphology rules that allow the system to produce automatically all the forms of inflected words from the base forms that the authors enter. Although the morphology rules cover the most common irregular forms, they may occasionally fail to generate the correct forms of words whose inflection is irregular, and this will be reflected in the previews of the generated texts. Another popup menu displays a table that allows the domain authors to correct those word forms by hand.

In a similar manner, new verbs can be added to the domain-specific lexicon, which can then be used with micro-planning expressions. As with nouns, verb entries are treated as triplets containing information about equivalent verbs for each language, and there are mechanisms to hide from the author features that require linguistic expertise and to generate automatically all verb forms. Morphology rules are particularly important in Greek and Italian, where verbs are highly inflected.

4.3 Micro-planning expressions

Micro-planning expressions (Figure 5, bottom right) specify how relationships and attributes can be conveyed in the supported languages; for example, which verb to use to express a relationship, how the participants of the relationship relate to the subject and object of the resulting phrase, etc. As discussed previously, relationships and attributes are represented in the authoring tool by fields.

There are two types of micro-planning expressions, clause-plans and templates, and the domain author can select between the two [Androutsopoulos *et al*. 2001]. In clause plans, the domain author specifies the verb to be used (from those available in the domain-specific lexicon), the voice and tense of the resulting clause, the preposition (if any) to be included between the verb and the object of the clause, any adverb (*e.g*., "frequently") to be included, and strings to be concatenated as adjuncts at the beginning or end of the clause (*e.g*. "In case of fire, ring the bell.").

Templates provide much stricter control over the surface form of the expression than do clause plans. They contain sequences of elements, which, after being instantiated, are simply concatenated to produce the final expression. Among other things, they are able to produce forms of expression that are not clauses. Each slot can be filled either by a particular string, by an expression that refers to the owner of the field, or by an expression that refers to the filler of the field. In the last two cases, it is also possible to specify that a particular type of referring expression (*e.g*., pronoun) should be generated.

4.4 Macronode authoring

To assist in developing the MacroNode aspect of generation, which is significantly more difficult than the wholesale inclusion of completely canned text, a separate authoring tool has been created to allow for macronodes creation. Each small piece of text must contain a conditional statement specifying when it will be included in the final text along with its placement in the macronode's message graph. This entails entering information on rhetorical type, pronominalization, user modelling parameters, etc. The MacroNode authoring tool allows these structures to be created by museum staff using a GUI instead of editing the XML structures that the system itself uses. Because the text fragments of the content units associated with each macronode are written as a fixed string, museum staff are able to create complete texts more quickly and easily than using the full-scale authoring system. This has benefits, such as rapid prototyping, which allows exhibit authors to decide whether particular exhibit descriptions are worth the cost of being more deeply developed. A previewing capability is also included, allowing the author to change different parameters such as user type, language and level of expertise, and then see the regenerated text to ensure it has been properly structured.

4.5 Reusing existing language resources

M-PIRO reuses a large-scale systemic grammar of English, which derives from ILEX and WAG [O'Donnell 1996], and the Italian and Greek grammars of M-PIRO reuse large parts of that grammar. The Upper Model, which has been used in several natural language generation applications, is another existing resource employed in M-PIRO. Monolingual large-scale lexicons in electronic form, such as WordNet and LDOCE [Fellbaum 1998; Wilks *et al*. 1996], constitute an alternative type of semantic/linguistic resource to M-PIRO's domain-dependent lexicon. As in most natural language generation applications, however, we expect M-PIRO's domain-dependent lexicon to be rather small. In most domains, it will typically contain a few dozens of lexical entries, many of which will be for specialized terms (*e.g*., "kouros", "stater" in the current domain) that are not included in general-purpose lexicons. Hence, the potential gain from using existing monolingual large-scale lexicons is rather small.

Multilingual domain-specific glossaries are of greater value, as they often provide exact mappings of terms across languages. They would be particularly useful when the speakers of the other languages that the domain author has access to are not domain experts, and hence may not be certain about the appropriate terms to be used in their languages. Even in those cases, however, using the glossaries is more a matter of offering easy access to them, perhaps by installing them on the same workstations as the authoring tool, rather than a matter of integrating them into M-PIRO's technology. Overall, then, the potential gain from attempting to incorporate existing lexical resources into M-PIRO's technology is rather small.

5. SPEECH SYNTHESIS

Another significant aspect of the M-PIRO project is the simultaneous generation of speech coordinated with the textual, web-based presentation. Adding audio allows M-PIRO to be used in auditory-only interfaces (e.g. mobile environments) or in applications where the text cannot be shown (e.g. a 3D virtual reality environment). Traditionally, the quality of synthetic speech is not very natural since generic models are used in order to cover documents from almost any domain (e.g. news, law, or culture). In M-PIRO we exploited the constraints of this specific domain in order to deliver more natural synthetic speech by accentuating focus information and rendering more realistic, intonational curves. To achieve this, we set up an enriched pipeline between the text generator and the text-to-speech system that transmits meta-information about the actual spoken text, thus allowing us to use more linguistic phenomena during prosody generation. This pipeline is based on the SOLE markup notation [Hitzeman *et al*. 1999] and has been extended according to the domain requirements. We have adapted the Festival speech synthesis system to improve on existing English speech and add support for Italian. Furthermore, we developed the DEMOSTHeNES Speech Composer in order to efficiently accommodate Greek resources and host Greek speech synthesis.

Traditional text-to-speech conversion fails to vocalize texts in a rhetorical way due to the generalizations of the underlying prosodic models and the lack of a reliable pragmatic analysis to detect the context of the document. Moreover, when going from text to speech, all the stages of this process introduce an error quantity: e.g. part of speech identification fails in 3% of the cases for Greek using statistical taggers. In addition, intonation events are often quantized to a set of 4-5 marks in order to avoid training from large data sets and thus reduce the probabilities of getting exceptional values in the F0 curve (e.g. ToBI annotation offers more than 20 discrete mark combinations, while only 4-5 of them are used in generic text-to-speech systems). In M-PIRO, such problems were addressed by taking into account the constraints of its domain and the ability of the text generator to provide the speech synthesizers with information that either enhances the performance of the above models or replaces them with highly accurate data.

5.1 Integration of text generator and speech synthesizers

The SOLE component lies between the text generator and the speech synthesizer. The SOLE markup language itself provides enumerated word lists and syntactic tree structures. On the syntactic tree, information exists at the phrase level about the phrase type (sentence, noun phrase, prepositional phrase, relative clause, etc) as well as at word level about the part-of-speech (determiner, noun, verb, preposition, etc.) (Figure 6). This structure carries error-free phrasing and POS information.

```
<utterance>
<relation name="Word" structure-type="list">
<wordlist>
…
<w id="w7">που</w>
<w id="w8">δημιουργήθηκε</w>
<w id="w9">κατά</w>
<w id="w10">τη</w>
<w id="w11">διάρκεια</w>
<w id="w12">της</w>
<w id="w13">αρχαϊκής</w>
<w id="w14" punct=".">περιόδου</w>
…
</wordlist>
</relation>
…
<elem phrase-type="S">
<elem lex-cat="PRP" href="words.xml#id(w7)"/>
<elem lex-cat="V" href="words.xml#id(w8)"/>
<elem phrase-type="PP">
<elem lex-cat="IN" href="words.xml#id(w9)..id(w11)"/>
<elem phrase-type="NP" newness="new" arg2="true" proper-group="true"
genitive-deixis="true">
<elem lex-cat="DT" href="words.xml#id(w12)"/>
<elem lex-cat="N" href="words.xml#id(w13)..id(w14)"/>
</elem>
</elem>
</elem>
…
</relation>
</utterance>
```

Figure 6: A SOLE example

Most of the sentences generated by the text generator in M-PIRO can be annotated with such detailed meta-information. Pieces of canned text integrated in the presentation are marked as "CANNED-TEXT" without any POS or phrasing information. Such cases are dealt with by the standard text-to-speech path, using instead trained models from the specific domain, thus leading to more natural speech.

The linguistic information carried by SOLE is being used to define better prosody. However, since the main aspect of prosody is to accentuate parts of emphasis in text, we concluded that the elements of the original specification of SOLE on their own could not always provide certain intonational focus points, even for the limited number of sentence structures generated for this domain.

Taking into account the capabilities of Exprimo, we extended the SOLE specification to accommodate elements that could directly or indirectly imply emphatic events for the specific domain. These elements stand for noun phrases and are:

- Newness information (to the particular user): `newness` [new/old]
- Number of times mentioned before (to the particular user): `mentioned-count` [integer]
- Whether they are a second argument to the verb: `arg2` [yes/no]
- Whether there is deixis: `genitive-deixis, accusative-deixis` [yes/no]

Furthermore, for the Greek language we also added to DEMOSTHeNES the ability to support rhetorical questions and explicit-emphasis. Some of these features are handled by trained models while the others are sent to an e-Text to Speech and Audio (e-TSA) component [Xydas and Kouroupetroglou, 2001b] to identify prosodic contours. The combination of the above elements provides the synthesizers with clues of stress and intonational focus.

5.2 Trained prosodic models

An important goal in M-PIRO's speech synthesizers is the delivery of more natural speech through improved prosody that approximates the prosody of human guides. To achieve that, professional speakers were used to record corpora and allow us to capture the spoken expressions of a museum guided tour. These corpora were created, one per language, by reading museum exhibit descriptions. For example, in the case of Greek, a 30000-word corpus of exhibit descriptions was created. The text corpus was first annotated by DEMOSTHeNES and was then produced in a readable format for the speakers to read them out following the annotation directions. The produced corpora were further segmented and annotated using ToBI marks. These marks were trained against the linguistic phenomena presented in the SOLE documents to produce the trained prosodic models, using Classification and Regression Trees [Breiman et al. 1984]. Training was split in 3 parts:

- Break prediction, where break indices were assigned to each word
- Accent prediction, where accent tones were assigned to stressed syllables
- Endtone prediction, where ending tones were given to phrase boundaries.

In most cases we used the Linear Regression model [Black and Hunt, 1996] for F0 rendering. These LR models were also trained from the above corpora and thus led to the production of smooth intonational curves for the specific domain.

5.3 Speech Synthesizers

The Festival language independent speech synthesis system in M-PIRO accommodates Italian and English. DEMOSTHeNES Speech Composer hosts Greek synthesis. Both systems implement Heterogeneous Relation Graphs [Taylor *et al.* 2001], to easily exploit XML documents (SOLE, VoiceXML, etc). Running in server mode, the systems are connected to the Speech Controller component, which, depending on the user's language preference, routes the text to the appropriate synthesizer. In all languages, diphone synthesis was used. Diphones were extracted from a dedicated voice database with isolated nonsense words.

The speech synthesis output quality is critical, especially in the Virtual Reality application where it is the only form of exhibit description since the use of text is not considered an appropriate metaphor for a 3D environment. The speech synthesis systems developed in M-PIRO were demonstrated at the Foundation of the Hellenic World's immersive projection-based system "Magic Screen" [Roussou 2000] where users, wearing special 3D shutter glasses, could select a virtual object and listen to its description and associations.

6. CONCLUSION

The M-PIRO project provides tools and techniques to enhance a museum visitor's experience by facilitating information delivery in a number of different languages and in a personalized way. It allows textual and spoken descriptions of exhibits to be generated automatically from an underlying language-neutral database and existing free-text descriptions. Each text is tailored according to the user's interests, background knowledge, and language skills.

A prototype that generates personalized descriptions of 50 archaeological exhibits in the form of web pages has been developed, which shows the potential of M-PIRO's technology in information presentations over the web and at public access kiosks. A second prototype delivers personalized audio descriptions of exhibits in an immersive virtual reality environment, using a tight integration between the text generator and the speech synthesizer to achieve more natural prosody. Authoring facilities allow domain experts to configure M-PIRO's technology for new exhibit collections, and to control both the information that is delivered to the visitors and the language expressions that are used to convey the information.

The M-PIRO consortium consists of: the University of Edinburgh (UK, co-ordinator), ITC-irst (Italy), NCSR "Demokritos" (Greece), the University of Athens (Greece), the Foundation of the Hellenic World (Greece), and System Simulation Ltd (UK). For more information, consult: http://www.ltg.ed.ac.uk/mpiro/.

7. REFERENCES

Androutsopoulos, I., V. Kokkinaki and A. Dimitromanolaki. "Dynamic Multilingual Object Descriptions: The M-PIRO Project", presented at the 29th Conference on Computer Applications and Quantitative Methods in Archaeology, Gotland, Sweden, April 25-29, 2001

Bateman, J.A. "Enabling technology for mulilingual natural language generation", *in Natural Language Engineering*, 3(1):15 55, 1997

Bateman, J.A.. "The theoretical status of ontologies in natural language processing", in S. Preuss and B. Schmitz, editors, *Text Representation and Domain Modelling – Ideas from Linguistics and A*I, pages 50–99. Technische Universitat Berlin, KIT-Report 97. 1992

Bateman, J., C. Matthiessen, C. Nanri and L. Zeng. "The re-use of linguistic resources across languages in multilingual generation components", in *Proceedings of IJCAI'91*, Sydney, Australia, 1991

Black, A. and Hunt, A. "Generating F0 contours from ToBI labels using linear regression", *ICSLP96*, vol. 3, pp 1385-1388, Philadelphia, PA. 1996

Breiman, L., Friedman, J. Olshen, R. and Stone, C. *Classification and regression trees*, Wadsworth and Brooks, Pacific Grove, CA. 1984

Callaway, C., Daniel, B. and Lester, J. "Multilingual Natural Language Generation for 3D Learning Environments", in *Proceedings of the 1999 Argentine Symposium on Artificial Intelligence*. Buenos Aires, Argentina, 177-190, 1999

Cosi, P., F. Tesser, R. Gretter and C. Avesani. "Festival Speaks Italian!", in *Proceedings of the 7th European Conference on Speech Communication and Technology*, EUROSPEECH 2001, Aalborg, Denmark, September 2001

Dimitromanolaki, A., I. Androutsopoulos, V. Karkaletsis, and C.D. Spyropoulos. "The Greek generation component". Deliverable 1.3, M-PIRO, 2001

Dimitromanolaki, A., I. Androutsopoulos and V. Karaletsis. "A Large-Scale Systemic Functional Grammar of Greek", presented at the Int. Conference on Greek Linguistics, Sorbonne, France, 2001

Fellbaum, C. (editor). *WordNet – An Electronic Lexical Databas*e. MIT Press, 1998

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

Hitzeman, J., A. Black, C. Mellish, J. Oberlander, M. Poesio and P. Taylor. "An annotation scheme for concept-to-speech synthesis", in *Proceedings of the European Workshop on Natural Language Generation*, pp. 59-66, 1999

Isard, A., J. Oberlander, I. Androutsopoulos and C. Matheson. "Speaking the Users' Languages". *IEEE Intelligent Systems*, vol. 18, no. 1, pp. 40-45, 2003

Mann, William C. and Sandra A. Thompson. *Rhetorical structure theory: A theory of text organization*. Technical Report ISI/RS-87-190, USC Information Sciences Institute. 1987

Melengoglou, Alexander. *Multilingual Aggregation in the M-PIRO System*. Master's thesis, University of Edinburgh. 2002

Mellish, C., M. O'Donnell, J. Oberlander, and A. Knott. "An architecture for opportunistic text generation", in Proceedings *of the 9th INLG Workshop*. Niagara-on-the-Lake, Ontario. 1998

Not, E., and M. Zancanaro. "Building Adaptive Information Presentations from Existing Information Repositories", in *Proceedings of the International Workshop on Information Presentation and Multimodal Dialogue*, Verona, Italy, 14-15 December, 2001

Not, E., and M. Zancanaro. "The MacroNode Approach: mediating between adaptive and dynamic hypermedia", in *Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems*, AH'2000, Trento, August 2000

Not, Elena, Ion Androutsopoulos, Jo Calder, Aggeliki Dimitromanolaki, Vangelis Karkaletsis, and George Paliouras. "Generation architecture and description of first M-PIRO proto-type". M-PIRO Deliverable 1.2, University of Edinburgh. 2001

Oberlander, J., M.O'Donnell, A. Knott and C. Mellish. "Conversation in the Museum: Experiments in Dynamic Hypermedia with the Intelligent Labelling Explorer", in *New Review of Hypermedia and Multimedia*. 4:11-32, 1998

O'Donnell, M., C. Mellish, J. Oberlander, A. Knott. "ILEX: An Architecture for a Dynamic Hypertext Generation System", in *Natural Language Engineering*, 7(3): 225-250, 2001

O'Donnell, M. "Input specification in the WAG sentence generation system", in *Proceedings of the 8th International Workshop on Natural Language Generation*, Herstmonceux Castle, UK, June 1996

Paris, C., K. Vander Linden, M. Fischer, A. Hartley, L. Pemberton, R. Power, and D. Scott. "A support tool for writing multilingual instructions", in *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1398 1404, Montreal, Canada, 1996

Power, R., and N. Cavallotto. "Multilingual generation of administrative forms", in *Proceedings of the 8th International Workshop on NLG*, pages 17-19, Herstmonceux Castle, UK, 1996

Power, R., and D. Scott. "Multilingual authoring using feedback texts", in *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and the 17th International Conference on Computational Linguistics* (COLING-ACL 98), pages 1053-1059, Montreal, Canada, 1998

Reiter, Ehud. "Has a Consensus NL Generation Architecture Appeared, and is it Psycholinguistically Plausible?", in *Proceedings of the Seventh International Workshop on Natural Language Generation. Kennebunkport*, ME, 163-170, 1994

Reiter, Ehud. "NLG vs. Templates", in *Proceedings of the Fifth European Workshop on Natural Language Generation*, The Netherlands, 1995

Reiter, Ehud and Robert Dale. *Building Natural Language Generation Systems*. Cambridge, U.K.: Cambridge University Press. 2000

Roussou, M. "Virtual Heritage: From the Research Lab to the Broad Public". In Franco Niccolucci (ed.), *Virtual Archaeology Proceedings of the VAST2000 Euroconference*, Arezzo, Italy. Oxford, Archaeopress, 93-100, 2002

Shaw, J. and K. McKeown. "An architecture for aggregation in text generation", in *Proceedings of the 16th International Joint Conference on Artificial Intelligence* (IJCAI'97). Nagoya, Japan, 1997

Taylor, P., Black, A., and Caley, R. "Heterogeneous Relation Graphs as a Mechanism for Representing Linguistic Information", *Speech Communications 33*, pp 153-174, 2001

Taylor, P., A.W. Black and R. Caley. "The Architecture of the Festival Speech Synthesis", in *Proceedings of ESCA 98 Workshop*, 1998

Wilks, Y., B.M. Slator, and L.M. Guthrie. *Electric Words – Dictionaries, Computers, and Meaning*s. MIT Press, 1996

Xydas, G. and Kouroupetroglou, G. "DEMOSTHeNES Speech Composer", in *Proceedings of the 4th ISCA Tutorial and Workshop on Speech Synthesis*, SSW4, pp. 167-172, Perthshire, Scotland, 2001

Xydas, G., and G. Kouroupetroglou. "Text-to-Speech Scripting Interface for Appropriate Vocalization of e-Texts", in *Proceedings of the 7th European Conference on Speech Communication and Technology*, EUROSPEECH 2001, Aalborg, Denmark, September 2001