Procesamiento del Lenguaje Natural, Revista nº 40, marzo de 2008, pp. 17-26 recibido 28-01-08, aceptado 03-03-08

Aggregation in the In–Home Domain*

Agregación en el entorno domótico

Eva Florencio, Gabriel Amores, Guillermo Pérez, Pilar Manchón
Grupo de Investigación *Julietta*
Universidad de Sevilla
Palos de la Frontera, s/n
41004 Sevilla, Spain
{evaflorencio,jgabriel,gperez,pmanchon}@us.es

Resumen: Este artículo describe experimentos realizados con vistas a determinar las preferencias de agregación léxica y sintáctica en inglés y español. El objetivo final es la implementación de dichas estrategias en el módulo de generación de lenguaje natural de un sistema de diálogo multimodal para el entorno domótico.
Palabras clave: Agregación, Generación de Lenguaje Natural, Sistemas de Diálogo

Abstract: This paper describes experiments carried out in order to determine syntactic and lexical aggregation preferences by English and Spanish users. The final goal of this work is the implementation of such strategies in the NLG module of a multimodal dialogue system in the in–home domain.
Keywords: Aggregation, Natural Language Generation, Dialogue Systems

1 Introduction

Describing the state of the different devices in a scenario such as the one in Figure 1, where information can be presented and expressed in multiple ways, involves a great complexity for *Natural Language Generation* (NLG) systems, and even for human beings.

Figure 1: Virtual House Example

Thus, the house in Figure 1, could be described by focussing on those devices which are switched–on, or we could group them according to their location, or type, as shown in examples 1, 2 and 3, respectively:

(1) The TV, the lights in the sitting room and the light in the kitchen are on.

(2) In the sitting room, the light is on. The light is on in the kitchen and the TV is on in the bedroom.

(3) The lights in the sitting room and kitchen are on, and the TV in the bedroom is on.

Moreover, not only can elements be grouped in several ways, but information can also be aggregated differently. Thus, the state of each individual device could be described by single independent clauses without combining them, as shown in example 4:

(4) The light in the bedroom is off. The blinds in the bedroom are rolled down. The TV in the bedroom is off. The lights in the patio are off ...

Although this way of presenting information is perfectly grammatical, it results in very monotonous and machine–like outputs. An NLG system which is capable of performing different aggregation strategies will produce a more natural output.

This paper describes experiments carried out in order to determine aggregation preferences by English and Spanish users. The final goal of this work is the implementation of such strategies in the NLG module of a

* This work has been funded by the Education and Science Spanish Ministry under the project *GILDA: Natural Language Generation for Dialogue Systems* (TIN2006-14433-C02-02).

multimodal dialogue system in the in–home domain.

The paper is organised as follows. Section 2 introduces the process of aggregation and its relevance in natural language generation. Next, section 3 describes the MIMUS multimodal dialogue system in which the aggregation strategies will be implemented. Section 4 outlines the initial working hypothesis to be confirmed by the experimental results. The experiments carried out are described in section 5. Sections 6 and 7 review the results obtained and the conclusions to be drawn from the experiments. Finally, section 8 advances some of the lines to be carried out from this moment in the context of the project.

2 Aggregation

A review of the literature on aggregation (Dalianis, 1999; Wilkinson, 1995; Shaw, 1998; Cheng, 2000) clearly points out that there is no agreement on its definition or where to place it in the generation process. Albeit thorough attempts have been made to come up with a core definition (Reape and Mellish, 1999) and a standard architecture (Cahill and Reape, 1999), conceptual problems arise.

For the purpose of this project, aggregation is conceived of as a process which removes redundant information from a text because it can be inferred or retrieved from linguistic sources (the remaining text), from computational sources (ontology), or pragmatically (using common knowledge).

In this work, we will focus on syntactic aggregation, understanding it as the process of combining phrases by means of syntactic rules, such as coordination, ellipsis or subordination. There are, however, some cases of lexical aggregation covered in this study too. Lexical aggregation is understood as the process of mapping several lexical predicates/lexemes into fewer lexical predicates/lexemes.

Pronominalisation is considered as a special case of lexical aggregation on the basis of Quirk et al. (1985)'s analysis of pro–form reduction. The theoretical motivation for it is that, indeed, it reduces the number of lexemes or predicates, but it is done by means of a pronoun, unlike other cases of reduction.

We claim that all these phenomena have a linguistic motivation and, consequently, they should be linguistically–grounded. As noted by Reape and Mellish (1999), most NLG systems lack a linguistic foundation to account for aggregation strategies.

3 The MIMUS Dialogue System

The context for this project is MIMUS, a multimodal and multilingual dialogue system based on the Information State Update (ISU) Approach (Larsson and Traum, 2000). The system has a symmetric architecture that allows that both the input and the output can be presented in graphical, voice or mixed (voice plus graphical) modalities. Besides, as it is a multilingual system, the user may interact dynamically in English and Spanish (Solar et al., 2007). MIMUS is made up of a series of collaborative agents (Pérez, Amores, and Manchón, 2006) that cooperate and communicate among them under the Open Agent Architecture (OAA, Martin, Cheyer, and Moran (1999)) framework.

The core module is the Dialogue Manager (DM), a collaborative agent that is linked to a Natural Language Understanding (NLU) module and to a Generation Module. Dialogues are driven both by the semantic information provided by the user and by the dialogue expectations generated by the dialogue manager. MIMUS incorporates its own specification language for dialogue structures that allows for the representation of the dialogue history, the control of expectations and the treatment of ambiguity.

The current version of MIMUS contains a hybrid NLG module in which sentence planning takes the form of predefined templates, as described in (Amores, Pérez, and Manchón, 2006). Utterances are elaborated from the mapping of abstract content representations to linguistic ones. In addition, some canned texts are used for common invariable expressions such as *Hello*, *Thank you*, or *Bye–bye*.

4 Working Hypothesis

The final goal of this work is to implement aggregation strategies in our NLG system. Namely, the final NLG module will be required to produce coordinated messages as well as sentences containing other linguistic phenomena, such as ellipsis, gapping or stripping. For instance, sentence 5 below shows an example of how the system should be able to concatenate the light's locations, either by juxtaposition or coordination, and produce

ellipsis or contribute with cue words such as *also*.

(5) The lights are on in the sitting room, in the bedroom, and in the kitchen. The hall is also on.

4.1 Location in the overall system

This section discusses where aggregation strategies could be placed in the NLG module of MIMUS.

Our first hypothesis is that both syntactic and lexical aggregation in the generation process in MIMUS will be located in the sentence planner. That is, sentence planning templates will be expanded with linguistic information so that they can perform syntactic and lexical aggregation.

As explained in the previous section, sentence planning templates map conceptual representations into linguistic ones that will later be passed on to the surface realiser. Therefore, the type of syntactic construction should be specified in the sentence planner so that the surface realiser transforms it into a linguistic unit by means of syntactic rules. The form that terminal nodes will have if lexical aggregation has taken place should also be specified. For instance, some items may have been lexically aggregated by employing a hypernym (e.g., *device*) instead of their hyponyms (e.g., *light*, *TV*, *fan* and/or *blind*).

In this fashion, the proposed architecture including aggregation can be seen in Figure 2.

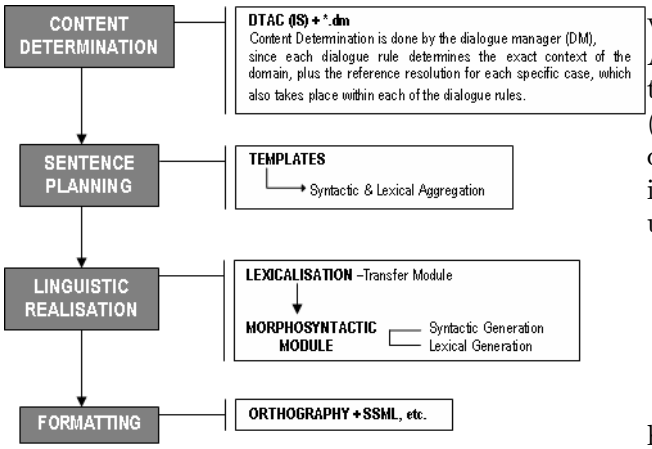

Figure 2: Proposed location of aggregation strategies in the NLG module

4.2 Linguistic constructions expected

With a view to implementing aggregation in the NLG module of our system, it is important to have some understanding of the grammatical coverage needed in the in–home domain. In addition, the linguistic coverage of the expected texts to be generated is also conditioned by the type of application being implemented (a multimodal dialogue system), and the type of interactions supported (requests about the state of devices in the in–home domain).

Taking into account possible questions that users may formulate when interacting with the system, answers may reply to questions about:

a. Quantity: the number of device(s) satisfying a specified condition(s).

b. State: the state (on or off) of the devices will be requested. Two subtypes may be found:

- Replies about the state of devices (*How is the light in the kitchen?*)
- Confirm the state of devices (*Is the light in the kitchen on?*)

c. Devices: information about which devices are in a specific state or location, i.e. (*Which devices are on in the house?*)

d. Location: obtain information about the location of devices, i.e. *Where is the tv?*

As discussed in Section 1, the information gathered may be grouped according to some common feature, for example, the type of device, the state they are in, or the location. As a first hypothesis, our prediction is that the grouping will mainly be done by location (see example 6 below), perhaps as a consequence of the distribution of the house, which is clearly separated into rooms, as seen in Figure 1.

(6) In the sitting room, the light is on, the fan is off, and the TV is on. In the bedroom, all the devices are on. In the patio, one light is on.

Nevertheless, the description could also hinge on the type of device or on their state. In those situations in which one of these characteristics (state, device or location) is explicitly mentioned in the question, it is foreseen that:

1. If the device is explicitly mentioned, then the grouping is done by location;

   Sys: *Please, tell me the state of the lights.*
   Usr: *In the sitting room, there is one light on. In the hall, the light is on. In the kitchen, the light is off. In the bathroom, it is on. In the patio, two lights are on and two are off.*

2. If the location is explicitly mentioned, then the grouping is done by device type:

   Sys: *How are the devices in the sitting room?*
   Usr: *There is one light on and the other one is off; the TV is on and the fan is off.*

3. If the state is the only feature mentioned, then it is considered as a non–specific situation in which the general prediction applies (i.e., grouping will be done by location).

   Sys: *Which devices are on?*
   Usr: *In the sitting room, only the fan is on. In the bedroom, the light and the TV are on. In the hall, two lights are on.*

4.3 Types of aggregation required

Concerning the types of syntactic and lexical aggregation that will be necessary in the MIMUS dialogue system, what follows is a list of the ones that should be implemented. The system should be able to produce them, but also to combine them when necessary. Besides, the insertion of some cue words or discourse markers would also be desirable.

4.3.1 Syntactic Aggregation

The next syntactic aggregation processes are required:

- **Paratactic constructions:** linking units of the same rank (sentences, clauses or phrases –the latter case will be referred to as constituent coordination). They are used whenever we need to go through a list of references.

  – Coordination: *[The light in the kitchen is on] and [the blind is rolled up].*

  – Constituent coordination: *[[The light in the kitchen] and [the light in the garage]] are on.*

- **Reduction:** It is probably the most common definition of aggregation in the literature and one of the most controversial aspects of its definition. Reduction is the process of removing information that can be inferred or retrieved from the remaining text. Different kinds are distinguished, depending on the type of information elided.

  – Ellipsis: In our domain, we expect it to be performed mainly when asking about a particular device or when there is only one type of device in a location.

    (7) *The (light in the) patio is on.*

  – Gapping: It is prone to happen when the main verb is understood, because it has just been mentioned, or when it is a copulative verb. In this domain, the main verb will be the copulative *estar/to be* in almost every sentence.

    (8) *In the sitting room, the TV is on and the fan (is) off.*

  – Stripping: It will take place when describing a device that shares the same state as the one previously mentioned.

    (9) *The light is off and the stove [is off] too.*

- **Multiple aggregation:** more than aggregation process, including also lexical aggregation takes place. For instance,

  (10) *In the patio, there are two lights on and* [CONSTITUENT COOR] *one* [PRONOMINALISATION: light] *off. The* [ELLIPSIS: light in] *kitchen is on and* [COOR] *the bathroom* [GAPPING: is] *off.*

4.3.2 Lexical Aggregation

Reducing the number of lexemes or predicates is required when all the devices in the same location have the same state, for instance: *En el dormitorio, todo está apagado/In the bedroom, everything is off*; or when describing the same device, such as *Hay una luz encendida en el baño y otra en la*

cocina/*There is one light on in the bathroom and another one in the kitchen.*

Apart from these pronominalisations, we also expect users to make use of other types of lexical aggregation such as the use of hypernym instead of its hyponyms, as in *The devices are on* (instead of *The light and hob in the kitchen are on*)/*Los aparatos están encendidos* (instead of *La luz y la vitrocerámica están encendidas en la cocina*).

4.3.3 Cue Words

Finally, the following cue words may contribute fluency, cohesion and coherence to the output messages: *también*; *así como*; *tanto... como...*; and *sin embargo*, *salvo*, or *pero* in Spanish; and *too*, *also*, *both*, and *but* or *however* in English. This will also result in more varied and less repetitive sentences.

5 Experiments

This section describes the experiments carried out in order to corroborate or refute the working hypotheses.

5.1 Goals

The main goal of these experiments has been the study of syntactic and lexical aggregation in the in–home domain, both in English and Spanish. Experiments were carried out in both languages in order to determine, in the first place, if they differ in the way information is aggregated.

In doing so, aggregation *per se* will be studied (how do speakers aggregate?, how often?, in which order?) with the aim of obtaining a pattern which may serve as a model of behaviour for its subsequent implementation in the system.

5.2 Design

The experiment consisted in showing the informants fifteen print screens of the house in which the devices were in different state configurations. Informants were then asked to describe the state of the devices.

The questions to be answered were in the range of possible requests that users can formulate to the system in the real application. Our final goal is to achieve a natural, human–like, virtual butler for the house.

The scenarios were distributed as follows:

- 3 scenarios asked about **quantity**.
- 1 about **location**.
- 4 about **devices**.
- 2 about **devices** and **location**.
- 3 about **description**.
- 2 asked for confirmation of **state**.

The user's profile was not specific; the only feature they had in common was that they were naïve, in the sense that they did not have any previous knowledge of the overall functioning of the system. The role of the users was to describe what they saw in a natural manner. In other words, they had to reply as information came to their minds, without elaborating the utterances beforehand.

They were provided with some information prior to the experiments, such as the type of devices they may come across (lights, televisions ...) as well as the state they may be in (on, off ...) and the number of them in each location.

There are nineteen devices available in the house, distributed as follows:

Sitting room: two lights, a TV, a fan and a blind.

Bedroom: one light, a TV, a blind and a fan.

Kitchen: one light and the ceramic hob.

Bathroom: one light.

Garage: two lights.

Patio: four lights.

Hall: one light.

The first settings were considered as an initial contact with the system, in which only basic information could be obtained, being aggregation either basic or non–existent at all.

As the experiment moves on, the difficulty increases. Different states with different devices and locations are combined together to see how the user aggregates information:

- simple enumeration,
- use of cue words, and
- preferences either by location, type of device or state.

5.3 Corpus

The corpus of study was obtained after interviewing twenty–four informants, twelve in Spanish and twelve in English. As aforementioned, since no specific user profile was sought, informants do not share the same characteristics in both languages. Since each informant was presented with 15 print screens, a corpus of 180 descriptions has been obtained for each language.

5.3.1 Spanish Corpus

In the Spanish version of the experiment, twelve users were enrolled. Out of these twelve informants, only four were women; the rest were men. All of them were native speakers of Spanish. Their education level was high, meaning that except for one of the informants, all of them held at least a university degree (Master's Degree, PhD students and PhDs were also interviewed). Their age ranged between 25 and 44 years old. The average age was 27.1, the median was 26, the mode was 25, and the standard deviation was 5.51.

5.3.2 English Corpus

For the English version, another twelve informants were recruited. As opposed to the Spanish version, the majority of the users were women, there were only four men involved in the experiment.[1] Two of these informants were bilingual (one English and French, and the other Tamil and English), but both reside in English-speaking countries. The average education level was degree studies. Except for three users (two Master's Degree and Degree), the rest of them were college students.

The range of age was from 20 to 62 years old. The average age was 24.3, the median was 21.5, the mode was 20, and the standard deviation was 11.7. The informants' age distribution of both languages can be seen in Figure 3.

6 Results

In order to properly analyse the results, we first specified the kind of question being asked. That is, among the questions asking for quantity, for example, we broke down the

[1]The data survey collection was carried out to determine if personal aspects, such as age, sex, or cultural level, could have an influence on their answers. Since no differences were found, no further comment will be made on these aspects.

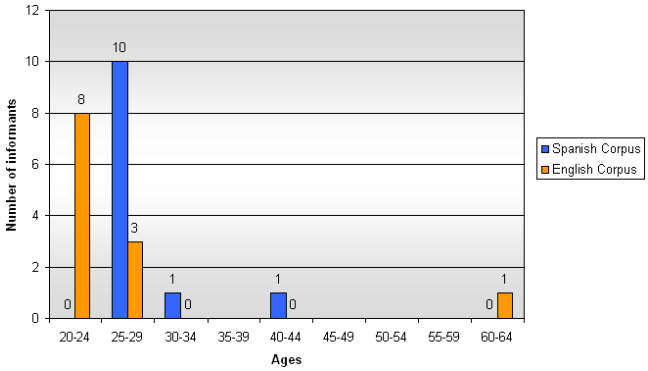

Figure 3: Users' age range in years

type of information demanded, determining if users were asked about the number of a specific device with a concrete state or about the number of devices in general, among other possibilities. Then, the different model answers were set and the usage percentages (out of the total answers for that specific kind of question) were given (see (Florencio, 2007) for further details).

At the same time, we also analysed the way in which informants grouped information, either by devices, states or location. After that, the lexical and syntactic aggregation found in each of the predominant patterns is pointed out, as well as the cue words used.

6.1 Spanish Results

6.1.1 Types of Syntactic and Lexical Aggregation Performed

The most common syntactic structures employed in Spanish were ellipsis, gapping and coordination (including constituent coordination), which were found in almost every reply. Coordination is the most frequent aggregation strategy employed (147 times), above all, when enumerating. Besides, since there were many questions demanding a description, it took place in almost every reply at least once (either sentence coordination or constituent coordination).

Ellipsis was the second most frequent type of aggregation (104 times), which was mostly used when the question specified the device. In such cases, most users elided the device in the reply.

Sys: *¿Qué luces están apagadas?* (setting 3)
Usr: *Las del salón, una del garage, la cocina, el baño, dos del patio y el dormitorio.*

Ellipsis also occurred when describing the state of a particular device.

Sys: *Dígame qué luces están encendidas.* (setting 6)
Usr: *Una (luz) en el salón,*
Una (luz) en el dormitorio,
dos (luces) en el garaje.

As expected, users avoided repetition when they deemed the information was inferable. Gapping was also used very frequently (81 times). There were some informants who omitted the main verb in 90% of their productions. This pattern was used by a few users regularly but not very often by the rest. The reason may reside in the copulative nature of the verb *estar*.

Sys: *¿Me puede describir el estado de todos los dispositivos (luces, aparatos y persianas)?* (setting 5)
Usr: *En el salón, las dos luces apagadas, televisión apagada, y ventilador en movimiento,*
la persiana del salón bajada,
la luz de la entrada apagada.
Las dos luces del garaje apagadas.
La luz de la cocina encendida,
la vitrocerámica encendida ...

Stripping was not used very frequently, with the exception of a couple of users who performed it (an average of twice per user, 4 times used). When used, it occurred when a location had more than one device, especially two, and both of them were in the same state, for example: *La luz de la cocina está encendida y la vitrocerámica (está encendida) también.*

Concerning lexical aggregation, *todo/a*, *ninguno/a*, *nada* (15 times), and *otro/a* (16 times) were often used when describing the same state or when all the devices shared the same state. *Otro/a* was often employed when enumerating the same device in different locations. No use of the hypernym *dispositivo(s)*, for instance, was made to refer to all lights, blinds, and so on; instead, *todo/ninguno* was preferred.

6.1.2 Use of cue words

The most commonly used cue word was *también* (15 times), in an average of at least one time per user. It was mostly used in enumeration. Some users alternated it with other cue words such as *así como* (1 time) or *tanto... como...* (2 times). Other markers used were adversative conjunctions, such as *sin embargo* (1 time), *pero* (1 time), *salvo* (1 time), and some distributive ones: *uno... otro...* (10 times). The words *sólo* and *el resto* were used once each.

6.2 English Results

6.2.1 Types of Syntactic and Lexical Aggregation Performed

An analysis of the syntactic and lexical aggregations performed on the English productions was carried out. With respect to syntactic aggregation, the most frequent strategies were ellipsis and coordination again.

Coordination, both sentence and constituent coordination, was employed in almost every utterance, adding to a total of 151 times. This phenomenon was employed when listing the types of devices and/or their locations. In the settings in which a description was required, coordination was mostly found.

Concerning reduction, ellipsis was highly employed as well. Ellipsis was realised 72 times in all. In the majority of cases, the type of device was the element elided in the sentence, particularly when it appeared in the question in hand. Another form of reduction used was gapping, which appeared 10 times. Only a couple of informants generally omitted the main verb in the sentence, even though it was a copulative verb. No other syntactic strategies were found.

With regard to lexical aggregation, we should point out the use of pronominal forms such as *one(s)* (16 times), *other/another* (5 times), *everything* (5 times) and *nothing* (2 times). They appeared mostly in descriptions, such as *Everything is off in the sitting room* or *The fan is off in the bedroom, but the one in the sitting is on.*

Finally, *all* (7 times) and *both* (15 times) were also employed in the descriptions when the same state applied to all the devices, either in the house or in a specific location: *All of the lights are on* or *Both of the lights are off in the sitting room.*

6.2.2 Use of cue words

It should be pointed out that English informants did not make use of many cue words in their replies. The most common cue words used were *also* (7 times) and the adversative *but* (9 times), which were used when enumerating or describing the state of all the devices in the house.

Other additive phrases employed were *as well as* (2 times), *so is...* (3 times), or *as is...* (1 time). For instance, *The light in the living room is on, so is the one in the patio*. As for other adversative phrases, the following ones were also mentioned: *except for* (1 time), *all the rest* (1 time), or *all the other* (3 times). An example would be *The light in the kitchen is on, all the rest are off*. The highly formal *as far as* was also used once when listing all the devices in the house (e.g. *As far as TV's, there are two*). The adverb *only* was employed just once to make a contrast, *On the patio, only one of the lights is off*.

7 Comparison and conclusions

By and large, the predictions and working hypotheses advanced in section 4.2, were mostly correct.

Grouping of information. With regard to the grouping of information, it was clearly done by location in both English and Spanish. This can be considered as a general preference on how to present the data as can be drawn form Figure 4.[2]

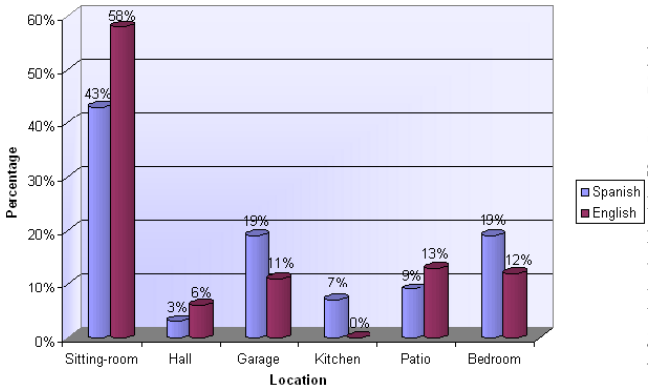

Figure 4: Preference for starting descriptions with location

Information was not only grouped by location, though; it was presented in a hierarchical way. This hierarchy was not the same for both languages. In Spanish, the most common way to present the data follows a [State — Device — Location] pattern (*Está encendida la luz de la cocina*); while, in English, the most popular pattern was [Device — State — Location] (*The light is on in the kitchen*).

[2]As we previously mentioned, this might be due to the graphical interface of the house.

Dialogue alignment. Another interesting result from the experiment was that sentence structuring in the replies aligned with the structure of the question. In both languages users were prone to reply following a similar pattern as the one employed in the question whenever a full sentence was provided. In both cases the end–weight and end–focus principles applied.

Long vs short answers. However, concerning the patterns established for the several questions, it should be highlighted that different models were obtained for English and Spanish. English speakers tend to construct full sentences, while Spanish speakers were more *economic*, and provided only the minimum information requested. For example, 53% of the Spanish informants replied to the quantity questions by just giving the number of devices, while only around 11% did so in English.

Another divergence is found in the patterns obtained for the reply location scenario. Nearly 70% of the Spanish users just provided the location, as opposed to a 75% of English speakers who provided full sentences (*The lights are on in the sitting room, in the bathroom, and in the hall*). This shows a preference for short incomplete sentences in Spanish and full sentences in English.

Syntactic aggregation. Another conclusion related to the preference for short or full sentences is the type of aggregation performed. As illustrated in Figure 5, Spanish users used more aggregation strategies than English informants, although not many aggregation strategies have been observed in the in–home domain overall. Apart from coordination, which was frequently employed in both languages, we could find other forms of syntactic aggregation in the Spanish corpus, such as ellipsis, gapping, and a few cases of stripping. Nevertheless, in the English data just ellipsis was found, and it was not commonly used. No other types of reduction were observed.

Lexical aggregation. As far as lexical aggregation, the results were very similar in English and Spanish. Pronominalisation was the most frequent strategy in both languages. We should emphasise the use of pronominalisation forms such as *todo/a*, *ninguno/a*, *nada*, *otro/a* in Spanish, and *one(s)*, *other/another*, *everything* or *nothing*

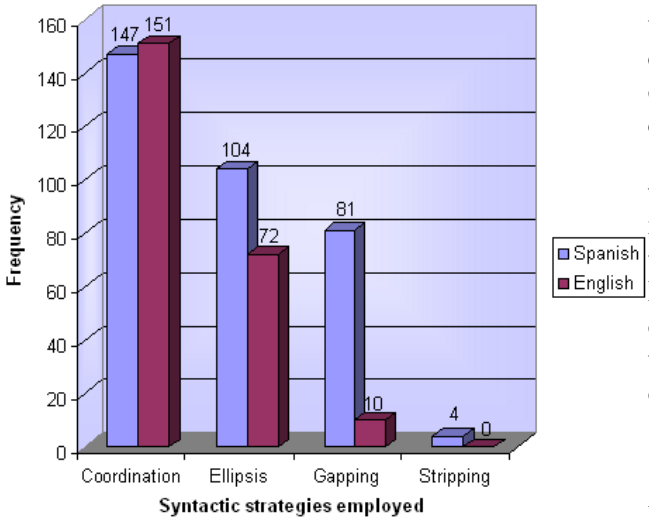

Figure 5: Syntactic aggregation performed

in English.

Use of cue words. Finally, with respect to cue words, no remarkable differences can be found between the two languages. *Also/también* obtained the highest frequency in both languages. The only point worth mentioning is that it seems that in English fewer cue words were employed but the ones employed were more varied. However, the difference is not significant.

Is aggregation language–dependent? Finally, although a much broader analysis should be performed, a comparison of the corpora in English and Spanish seems to suggest that aggregation is language–dependent instead of language–independent. Besides, the enormous differences found between the patterns established in each language plus the different aggregation strategies employed open the possibility of reconsidering the localisation of the aggregation process at a later stage (i.e., not in the Sentence Realiser, but on the Surface Realiser), or consider that the generation module as a whole should be language–dependent.

8 Future Work

At this point in the project, a new specification language is being created in collaboration with the TAP (a **T**ext **A**rranging **P**ipeline) project (Gervás, 2007) in an effort to create a set of interfaces which define generic functionality for a pipeline of tasks oriented towards natural language generation. The DTAC representation obtained by our dialogue system is currently being integrated with the TAP system so that different aggregation strategies for both languages can be compared on the basis of the results obtained by the experiments.

In addition, the new integrated prototype will incorporate preference strategies for lexical alignment, (i.e. if a user preferred the term *bombilla* instead of *luz* to refer to the lights in the house, the system should align consequently in the reply) and for fragmentary vs. verbose replies depending on the context.

References

Amores, G., G. Pérez, and P. Manchón. 2006. Reusing MT Components in Natural Language Generation for Dialogue Systems. *Procesamiento del Lenguaje Natural*, 37:215–221.

Cahill, L. and M. Reape. 1999. Component tasks in applied NLG Systems. Technical report, Information Technology Research Institute Technical Report Series.

Cheng, H. 2000. Experimenting with the Interaction between Aggregation and Text Structuring. In *Proceedings of the ANLP-NAACL 2000 Student Research Workshop*, pages 1–6, Seattle, Washington, USA.

Dalianis, H. 1999. Aggregation in Natural Language Generation. *Computational Intelligence*, 15(4):384–414, November.

Florencio, E. 2007. A study on syntactic and lexical aggregation in the in-home domain. Master's thesis, University of Seville, Spain, May.

Gervás, P. 2007. TAP: a Text Arranging Pipeline. Technical report, Natural Interaction based on Language Research Group, Facultad de Informática, Universidad Complutense de Madrid, May. Working draft.

Larsson, S. and D. Traum. 2000. Information state and dialogue management in the TRINDI Dialogue Move Engine Toolkit. *Natural Language Engineering*, 6(3-4):323–340.

Martin, D. L., A. J. Cheyer, and D. B. Moran. 1999. The Open Agent Architecture: A Framework for Building Distributed Software Systems. *Applied Artificial Intelligence*, 13(1-2):91–128.

Pérez, G., G. Amores, and P. Manchón. 2006. A Multimodal Architecture for Home Control by Disabled Users. In *Proceedings of IEEEACL Workshop on Spoken Language Technology (SLT)*, pages 134–137, Aruba, December.

Quirk, R., S. Greenbaum, G. Leech, and J. Svartvik. 1985. *A Comprehensive Grammar of the English Language*. Longman Group Limited.

Reape, M. and C. Mellish. 1999. Just what is aggregation anyway? In *Proceedings of the 7th European Workshop on Natural Language Processing*, pages 20–29, Toulouse (France), May.

Shaw, J.C. 1998. Clause Aggregation Using Linguistic Knowledge. In *Proceedings of the 9th International Workshop on Natural Language Generation*, pages 138–147, Niagara-on-the-Lake, Canada, August.

Solar, C. Del, G. Pérez, E. Florencio, D. Moral, G. Amores, and P. Manchón. 2007. Dynamic Language Change in MIMUS. In *Proceedings of the Eighth Interspeech Conference (INTERSPEECH 2007 Special Session: Multilingualism in Speech and Language Processing)*, pages 2141–2144, Antwerp, Belgium, August 27-31.

Wilkinson, J. 1995. Aggregation in natural language generation: Another look. Technical report, Co-op work term report, Department of Computer Science, University of Waterloo, September.