Using cohesive properties of text for Automatic Summarization

Maria Fuentes
Universitat de Girona (mfuentes@ima.udg.es)
Horacio Rodríguez
Universitat Politècnica de Catalunya (horacio@lsi.upc.es)

Abstract A system allowing extractive automatic summarization of textual documents is presented. The system is based on the cohesive properties of text, namely lexical chains, co-reference chains and named entity chains. In this way the system extend the well known lexical-chaining paradigm for summarization. The system has been applied to summarization tasks on Spanish agency news. Results of its evaluation and comparison with a couple of baseline systems are presented.

1 Introduction

We present in this paper an Automatic Summarization, AS, system that uses the cohesive properties of the text for selecting the most informative fragments for including in the summary. The system uses lexical chains, as indicative of lexical cohesiveness, as primary source for ranking segments of the text but uses as well complementary sources, namely co-reference chains and Named Entities, NE, chains. The system has been tested with a corpus of Spanish agency news and its results compared with another available summarization system.

The system we present here is an extractive informative summarization system based on the cohesive properties extracted from the text. The system aims to be language independent provided we dispose of the knowledge sources for carrying out the pre-process steps and the lexical chains ranking, basically the corresponding WordNet[1]. In the experiments reported here the system has been applied only to Spanish.

[1] Some lexical resource for measuring relatedness between lexical items is needed. In our experiments, for Spanish, we have used EuroWordNet (http://www.hum.uva.nl/~ewn/WordNet). For English, Priceton's WordNet can be used instead (http://www.cogsci.princeton.edu/~wn/w3wn.html).

Although based on cohesion, relevance of the lexical units, i.e. words, terms and NE, is incorporated in an indirect way through the corresponding chains.

Due to the characteristics of test summaries, the extractive unit we use is the paragraph. Some experiments using finer grained units, sentences and clauses, have been carried out but are not reported here. Several compression rates have been experimented although we report here only the corresponding to the test corpus.

The organization of the paper is as follows. We present first, after this introduction, a short review of the state of the art in AS. Then the three kind of cohesion and relevance indicators, lexical chains, co-reference chains and NE chains are presented and their importance in Summarization justified. Section 4 presents the overall architecture of the system. Section 5 deals with the empirical evaluation of the system and, finally, section 6 states some conclusions and current and future development of our work.

2 Some current trends in Automatic Summarization

AS has become in last years an active line of research, first promoted by TIPSTER's SUMMAC and more recently by the DUC[2] competition.

Initially reduced to textual, monolingual, single-document condensation task, AS has evolved for covering currently a wide spectrum of summarisation tasks (that can be classified along several dimensions: extracting vs. abstracting, indicative vs. informative, generic vs. query-based, background, vs. getting the news, restricted vs. unrestricted domain, textual vs. multimedia, single-document, SDS, vs. multiple-document, MDS) and applications (biographical summaries, medical patient summaries, e-mail, Web pages, news, support

[2] http://www-nlpir.nist.gov/projects/duc/2001.html

to IR relevance feedback, headlines extraction, meeting recording, ...).

A lot of different techniques have been applied to SDS for i) locating the relevant fragments, i.e. sentences, paragraphs, passages, of the document, ii) ranking these fragments by relevance and iii) producing the summary. Among them: using lexical chains, [6], [4], co-reference chains, [3], alignment techniques, [5], similarity and divergence measures, as MMR, [7], statistical models, as Bayesian models, [24], HMMs, Logistic Regression, [9], Machine Learning approaches, including decision trees and ILP, [15], [25], sentence reduction, [13], Information Extraction techniques, [14], topic detection-based systems, as [12], systems using the rhetorical structure of the document, [17]. Sometimes these techniques are combined, as in [16], [21], [26], [1].

When dealing with MDS new problems arise: lower compression factors implying a more aggressive condensation, anti-redundancy, temporal dimension, more challenging co-reference task, ... Clustering of similar documents plays now a central role. Selecting the most relevant fragments from each cluster and assuring coherence of the summaries coming from different documents are other important problems. Among the most important contributions to this issue we can find the reformulation approach of MULTIGEN, [19], [10], the use of Webclopedia in NEATS, [17], the centroid-based approach of MEAD, [23].

3 Cohesion and relevance indicators

Using the discourse structure of documents seems to be a good choice for single document summarization. Traditionally, two main components have been distinguished in the discursive structure of a source text: cohesion and coherence. Cohesion tries to account for relationships among the elements of a text, including reference, ellipsis, conjunction, and lexical cohesion. On the other hand, coherence is represented in terms of relations between text segments, such as elaboration, cause or explanation. Thus, coherence defines the macro-level semantic structure of a connected discourse, while cohesion creates connectedness in a non-structural manner.

We will focus in our system on cohesion features.

Lexical Chains (initially proposed in [20] and widely used for summarization and other NLP related tasks) try to identify cohesion links between parts of text by identifying relations holding between their words. Two pieces of text are considered to be lexically related not only if they use the same words, but also if they use semantically related words. This is a way to obtain a certain structure of a text based on the distribution of its content.

Identity chains are the most simple form of lexical chain. They are supposed to contain terms that refer to the same object. They are created by pronominal cohesion, lexical repetition or instantial equivalence and are always text-bound, because the relation of co-reference can be determined only in the context of a text. In contrast, similarity chains are not text-bound. Their elements are held together by semantic bonds. These bonds are supra-textual, with a language-wide validity. The two types of chains are important for cohesion analysis, however, the advantages of similarity chains over identity is that they can be computed without requiring deep text understanding. These lexical chains can be computed irrespective to the context in which related words actually occur.

Lexical Chains provide a representation of text that has been used for a variety of NLP tasks, including topic passage segmentation, detection of malapropisms, automated text summarisation or automatic hypertext generation. See [10] for details.

The general procedure for constructing lexical chains usually follows three steps:

1. Select a set of candidate words
2. For each candidate word, find an appropriate chain relying on a relatedness criterion among members of the chains. Usually relatedness of words is determined in terms of the distance between their occurrences and the shape of the path connecting them in WordNet.
3. If a chain is found, insert the word in the chain and update it accordingly.

Chains are scored according to a number of heuristics: their length, the kind of relation between their words, the position in the text where they start, etc.

One of the drawbacks of lexical chains is that they are insensitive to the non-lexical structure of texts, such as their rhetorical, argumentative or document structure. For example, they don't take into account the position of the elements of a chain within the argumentative line of the discourse, sometimes not even within the layout- or genre-determined structure of the document. Therefore, the relevance of chain elements is calculated irrespective of other discourse information and, consequently, the strength of lexical chains is exclusively based on lexic.

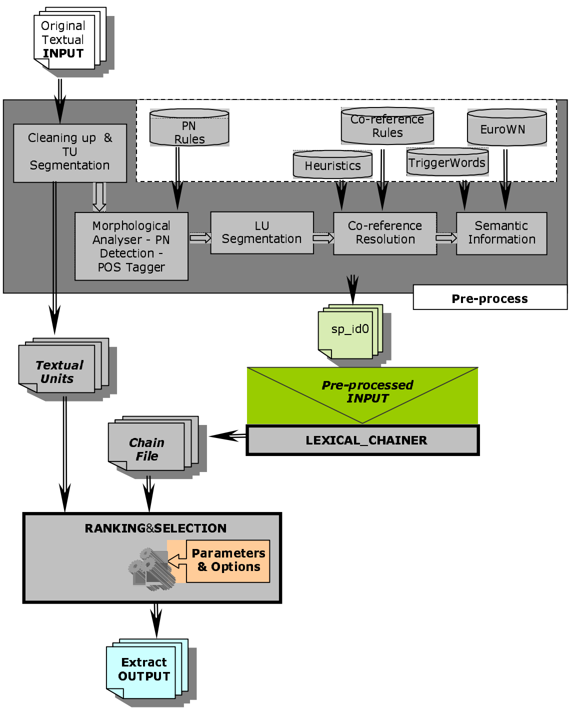

Figure 1 Architecture of the system

4 Architecture of the system

The overall architecture of the system is presented in Figure 1. As can be seen, the system performs in sequence three steps: 1) linguistic pre-process, 2) lexical chaining, 3) segment ranking and selection. Roughly, step 1 is in charge of segmenting the input document into textual and lexical units, extracting co-reference and NE chains and enriching the text with information needed for further steps, step 2 will identify lexical chains and score all three types of chains and step 3, finally, is in charge of scoring the textual units (segments) according to the scores of the chains traversing them and selecting the most appropriate segments for building the summary. We will address in turn these issues.

4.1 Linguistic pre-process

The first step of our system consists of cleaning up the input document and segments it in text units (TU).

Textual segmentation can be performed with varying degrees of granularity depending on the application, though in the experiments reported in this paper only simple paragraph segmentation is carried out[3].

The document is then processed by the standard morpho-syntactical analysis tools given by the CLiC-TALP system, [8].

This process includes morphological analysis, part of speech tagging and lexical unit (LU) segmentation. At the end the result consists of the original word, its lemma, and its part-of-speech.

The pre-process is completed by a NE recognition step, a co-reference step and a semantic tagging step.

NE are identified and classified using the system described in [2]. From detected NE, NE chains are build with a very simple string comparison mechanism. In the experiments reported here gazetteers have not been used because we wanted a neat comparison of summarisation using or not NE chains without the help of such resources. We may expect, however, that using accurate gazetteers will improve NER step and, thus, the quality of NE chains.

Co-reference links have been extracted only for some types of pronouns[4] using a simplified version of [22]. Because no parsing step has been included in our system we have dropped out all the constraints and rules involving syntactic information.

Semantic tagging has been performed attaching EuroWordNet, EWN, synsets to words, with *is-a* relations, and NE (with *instance* relations through the corresponding trigger word). No attempt to Word Sense Disambiguation has been made at this level. A

---

[3] The reason is that test summaries have been manualy extracted at paragraph level. See section 5.
[4] An extension for covering definite reference will be developed in the near future

fragment of a document after the pre-process is shown in figure 2.

```
1_1 México México NP00G00 tax: NP00G00:
geografía#4081235n
1_2 , , Fc *
1_3 23_may [23/5/??] W *
1_4 ( ( Fpa *
1_5 EFE EFE NP00O00 tax:NP00O00:organización
1_6 ) ) Fpt *
1_7 . . Fp *
1_8 - - Fg *
1_9 El el TDMS0 *
1_10 conservador conservador AQ0MS00 *
1_11 Vicente_Fox Vicente_Fox NP00SP0
tax: NP00SP0: persona#00004865n
1_12 , , Fc *
1_13 candidato candidato NCMS000 isa: 05840699n|
has_hyperonym:05840297n| has_hyperonym: 00004865n|
has_hyperonym: 00004473n|has_hyperonym:00002403n|
has_hyperonym: 00002728n| has_hyperonym:
00002403n| isa:05982191n|has_hyperonym: 06271584n|
has_hyperonym:05850058n| has_hyperonym: 00004865n|
has_hyperonym: 00004473n|has_hyperonym: 00002403n|
has_hyperonym:00002728n|has_hyperonym:00002403n|
1_14 del del SPCMS *
1_15 Partido_Acción_Nacional
Partido_Acción_Naciona NP00O00 tax: NP00O00:
organización
1_16 ( ( Fpa *
...
```

Figure 2 A fragment of a pre-processed document

4.2 Lexical chaining

The Lexical Chain system is the one proposed in [10]. It follows the work of [6] extending it for dealing with all three types of chains. Chain candidates are common nouns, proper nouns, named entities, definite noun phrases[5] and pronouns. For each candidate word, three kinds of relations[6] are considered:

- **Extra-strong**: Between a word and its repetition and between words belonging to the same synset.
- **Strong**: Between two words connected by a EWN relation.
- **Medium-strong**: If the link between the EWN synsets of the words has a length longer than 1.

In addition, this system establishes relations between common nouns and the rest of chain candidates, by means of the information provided by the semantic tagger. For instance, in figure 2 a relation *instance* between *Vicente_Fox* and a synset 00004865n, corresponding to *person*, is detected at pre-process step. This relation allows us to link a NE chain including *Vicente_Fox* and a lexical chain including any variant contained in the corresponding synset (ser humano, alma, persona, mortal, individuo, humano)[7].

To build the chains, there are constraints on the path length according the type of edges, determined by the relations between chain members. For the moment we only have considered only the extra-strong and the strong relations (mechanisms for taking profit of the rich relationship coverage of EWN have been implemented and initial experiments have been performed). For that reason our algorithm has a splitting chains process simpler than Barzilay's. When a new LU is processed by the system an attempt is made for attaching it to any of the existing lexical chains. In the case of identity chains the process is straightforward; in other cases different senses of the LU can be considered for being attached to different chains[8]. Anyway splitting is postponed as long as possible and limited as much as possible for preventing an uncontrolled growth of chains. A threshold mechanism for performing a pruning of less promising chains and maintaining the number of chains under control is used.

Once detected all the relevant chains of a document a process of scoring each chain is carried out. The score takes into account the length of the chain and its homogenousness. Once scored, strong chains are selected. A chain is considered strong if its score outperforms by twice the standard deviation the average of scores of all the chains. Only strong chains are considered for next step.

4.3 Segment ranking and selection.

TU are ranked and those crossed by most strong chains are considered to be most relevant. Another criterion is to consider the first TU crossed by a strong chain. This criterion favours the heading position of TU and obtains better results in the case of documents, as agency news, that frequently begins with a summary of the news.

Several forms of merging of the three kinds of chains could be considered. It is not a simple issue due to the different granularity of the involved units (word senses, definite phrases, named entities). For the moment only the

[5] Not implemented yet.
[6] Following Barzilay's nomenclature.
[7] human, soul, someone, person, mortal, individual
[8] "One sense for discourse" hypothesis is assumed.

simplest forms of combination are implemented, in the other cases chains are considered to be independent of each other and interact only when scoring the TU.

Once ranked, a certain number of TU is extracted from that list until a determined summary length is achieved.

4.4 Parameters of the system.

The following parameters can be defined in our system.

- Language: Catalan, English, Spanish and Multi-language.
- Directory where documents to be summarized are placed.
- File containing the identifiers of documents to be summarized.
- Directory where results have to be placed.
- Type of summary (monodoc, multidoc).
- Compression degree.
- Maximum distance in WN for medium-strong chains.
- Type of chain merging.
- Heuristic for scoring TU (1: first TU crossed by a strong chain, 2: TU crossed by maximum of strong chains).
- Relative scoring of 1st TU.
- Relative scoring of TU crossed by a strong chain.
- Relative scoring of TU crossed by NE chain.
- Relative scoring of TU crossed by a co-reference TU.

5 Empirical evaluation of the system

For evaluating our system a test corpus has been created within the framework of project Hermes (Hemerotecas Electrónicas, Recuperación Multilingüe y Extracción Semántica)[9]. The corpus consists of 120 news agency stories (reduced to 111 after removing news with only one paragraph) of various topics, including economy, finances, politics, science, education, sport, meteorology, health and society. Stories range from 2 to 28 sentences and from 28 to 734 words in length, with an average length of 275 words per story. The news were randomly selected from a corpus provided by EFE, the Spanish news agency.

From each news belonging to the set extractive summaries were manually built. 31 human evaluators were presented with several agency news articles. Each subject summarized a set of articles that went from 1 to 77. The objective was to have at less 5 different summaries made for each article.

The human summary was made via Web (figure 3 shows the main window of the evaluation page). Each news in turn was presented to the evaluator segmented at sentence level. Sentences were numbered so that they could be referenced easily. In order to deal with different compression degrees the human evaluators were asked to assign a score to each of the sentences of the article. Three possible scores, [0,1,2], were used to mark the relevance of the sentence in the whole article. In the instructions to the evaluators the term *relevance* was loosely defined. Essentially the meaning of relevancy 2 is “This sentence would occur in my summary” and the meaning of relevancy 0 is “This sentence wouldn’t occur in my summary”. Each evaluator was asked to provide as well for each document, a list key words.

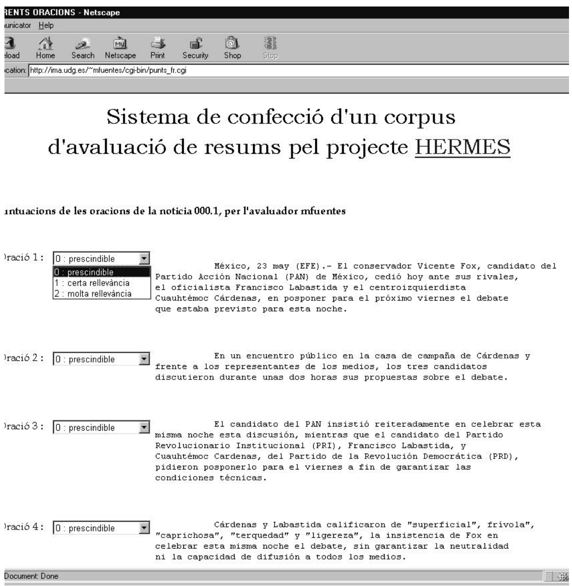

Figure 3 Interface for extracting summaries within Hermes project

[9] http://terral.ieec.uned.es/hermes

Two different golden standards were obtained from these scores, one containing summaries coming as close as possible to the 10% of the length of the original text (resulting on an average 19% compression) and the other containing the *best* summaries. We defined the *best* summary as a group of sentences with more than a half of the maximum possible score. This resulted on an average of 31% of the length of the original text (29% compression). Paragraph level extraction lead to better agreement between human evaluators and so this unit has been used for building both evaluation sets.

Using the *first* set of summaries as golden standard we have developed a set of experiments. The results are presented in table 1.

For the evaluation of our system against the golden standard we have used the evaluation software MEADeval[10] developed within the MEAD project. From this package we have selected the usual *Precision* and *Recall* measures and we have measured as well the *cosine*. Due to the characteristics of the documents of the corpus (agency news) the best summary can be built simply by selecting the first paragraph of the document. So, *cosine* measure could be in this case a more fair indicator of the goodness of the system.

Two baseline systems have been used for comparison: the *lead* method, i.e. extracting a number of paragraphs, starting on the first one, until the desired length, given the compression rate is achieved. The other baseline is *SweSum* system[11] a system allowing summarization of Spanish texts and a pretty way of customization.

Two heuristics schemata have been experimented (heuristic 1 and heuristic 2 in table 1). Heuristic 1 selects as most relevant the first TU crossed by a strong chain , heuristic 2 selects the TU crossed by maximum of strong chains.

First column in table 2 shows the main parameters governing the trial. *LexChains* means that lexical chains are taken into account, *PNChains* and *coRefChains* refer, respectively, to chains of NE and co-reference chains. *1stUT* means that a special weighting is assigned to the first TU.

[10] http://perun.si.umich.edu/clair/meadeval

[11] http://www.nada.kth.se/~xmartin/swesum/

As expected, given the characteristics of the documents baseline methods outperform our system except in the case of using heuristic 1 together with 1stUT. Our system presents however, in this case, a more balanced result of precision/recall/cosine figures (0.88, 0.88, 0.90). This is clearly a good indicator for documents not so biased towards leading summaries.

Another good indicator is the difference between precision/recall and cosine in the case of heuristic 2, that is less affected by leading effect. In the case of the two baselines there is a drop (from 0.95 to 0.90 and from 0.90 to 0.87), in the case of our method following heuristic 1 there is a small increment (0.82 to 0.85, 0.85 to 0.88, 0.83 to 0.87, 0.88 to 0.90) and in the case of our method following heuristic 2 the increment is larger (0.70 to 0.78, 0.73 to 0.81, 0.70 to 0.78, 0.82 to 0.86).

Regarding the influence of lexical chains, NE chains and co-reference chains on summarization, we can examine the results of Heuristic 2. Inclusion of NE chains has a positive (but small) effect on the accuracy. Inclusion of co-reference chains does not seem to affect the performance .

	PRECISION	RECALL	SIMPLE COSINE
Baseline			
Lead	0.95	0.85	0.90
SweSum	0.90	0.81	0.87
Heuristic 2			
LexChains	0.70	0.72	0.78
LexChains + PNChains	0.73	0.74	0.81
LexChains + PNChains + coRef Chains	0.70	0.71	0.78
LexChains + PNChains + coRef Chains + 1st UT	0.82	0.82	0.86
Heuristic 1			
LexChains	0.82	0.81	0.85
LexChains + PNChains	0.85	0.85	0.88
LexChains + PNChains + coRef Chains	0.83	0.83	0.87
LexChains + PNChains + coRef Chains + 1st UT	0.88	0.88	0.90

Table 1 Results of the experiments

Regarding NE chains our opinion is that the positive results could be improved with more accurate NE recognition/classification step. Using gazetteers will be a way of carrying off such improvement.

Regarding co-reference chains the main conclusion is that reduced to pronominal reference the resulting chains are simply too short for having a remarkable influence on the overall performance. Incorporating other forms of co-reference, specially noun phrase definite reference could be a good choice. Also developing more complex forms of chain merging could be a promising direction.

6 Conclusions and current and future developments

We have presented an informative extractive automatic summarization system. The system is based on the cohesive properties of text, namely lexical chains, co-reference chains and named entity chains. The system allows several forms of customization for experimenting different summarization schemata. Several experiments have been carried out. The results have been evaluated, compared with two baseline methods and discussed.

The system is being extended in several ways: i) improving the basic lexical chaining procedure for taking profit of rich semantic relations included in EWN, ii) incorporating more accurate NE recognition/classification modules, basically using gazetteers, iii) improving the, by now, rudimentary co-reference identificator, iv) experimenting more complex methods for merging the different chains and v) applying the system to other languages, specially English and Catalan.

References

[1] L. Alonso, M. Fuentes. 2002, Collaborating Discourse for Text Summarisation. In *Proceedings of the Seventh ESSLLI Student Session*, Trento

[2] M. Arévalo, X. Carreras, L. Màrquez, M. A. Martí, L.Padró, M. J. Simón. 2002, A Proposal for Wide-Coverage Spanish Named Entity Recognition. *Procesamiento del Lenguaje Natural*, Vol. 1, Num 3

[3] B. Baldwin, T. S. Morton. 1998, Dynamic Coreference-Based Summarization. In *Proceedings of the Third Conference on Empirical Methods in Natural Language Processing*. June 1998. Granada, Spain.

[4] M. Brunn, Y. Chali, C.J. Pinchak. Text Summarization Using Lexical Chains. In *Workshop on Text Summarization, ACM SIGIR Conference*. September 13-14, 2001, New Orleans, Louisiana USA.

[5] M. Banko, V. Mittal, M. Kantrowitz, J. Goldstein. 1999. Generating Extraction-Based Summaries from Hand-Written Summaries by Aligning Text Spans. In *Proceedings of PACLING-99* Waterloo, Ontario.

[6] R. Barzilay. 1997. *Lexical Chains for Summarization*, Master thesis. Ben-Gurion University of the Negev

[7] J. G. Carbonell, J. Goldstein. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *Proceedings of SIGIR 1998*

[8] Carmona, J.; Cervell, M.; Márquez, L.; Martí, M.A.; Padró, L.; Placer, R.; Rodríguez, H.; Taulé, M. & Turmo, J. 1998. An Environment for Morphosyntactic Processing of Unrestricted Spanish Text In *Proceedings of 1st International Conference on Language Resources and Evaluation, LREC'98*. Granada, Spain.

[9] J. M. Conroy, J. D. Schlesinger, D. P. O'Leary, M. E. Okurowski. 2001. Using HMM and Logistic Regression to Generate Extract Summaries. In *DUC Workshop on Text Summarization in conjunction with the ACM SIGIR Conference 2001,* New Orleans, Louisiana USA

[10] M. Fuentes. 2001. *Lexical Chains: Comparative study, implementation and application areas*, Technical Report, IMA 01-01-RR, Universitat de Girona

[11] V. Hatzivassiloglou, J. Klavans, M. Holcombe, R. Barzilay, M.Y. Kan, K.R. McKeown. 2001. SimFinder: A Flexible Clustering Tool for Summarization. In *NAACL'01 Automatic Summarization Workshop*

[12] E. Hovy, C-Y. Lin. 1999. *Automated Text Summarization in SUMMARIST* In *Advances in Automatic Text Summarization,* Mani and Maybury, eds.

[13] H. Jing. 2000. Sentence simplification in automatic text summarization. In *Proceedings of ANLP-2000*

[14] M-Y. Kan, K. Mc.Keown. 1999. *Information Extraction and Summarization: Domain Independence through Focus Types,* Columbia University Computer Science Technical Report CUCS-030-99

[15] K. Knight and D. Marcu. 2000. Statistics-Based Summarization - Step One: Sentence

Compression. In *17th National Conference of the American Association for Artificial Intelligence AAAI'2000*, Outstanding Paper Award, Austin, Texas, July 30-August 3, 2000.

[16] W. Kraaij, M. Spitters, M. van der Heijden. 2001. Combining a Mixture Language Model and Naive Bayes for Multi-Document Summarisation. In *Workshop on Text Summarization in conjunction with the ACM SIGIR Conference 2001*, New Orleans, Louisiana USA,

[17] C-Y Lin, E. Hovy. 2001. NEATS: A Multidocument Summarizer. In *Workshop on Text Summarization in conjunction with the ACM SIGIR Conference 2001*, New Orleans, Louisiana USA,

[18] D. Marcu. 1997. *The Rhetorical Parsing, Summarization, and Generation of Natural Language Texts* PhD Thesis, Department of Computer Science, University of Toronto,

[19] K. Mc.Keown, R. Barzilay, D.K. Evans, V. Hatzivassiloglou, M-Y. Kan, B. Schiffman, S. Teufel. 2001. Columbia Multi-Document Summarization Approach and Evaluation, (DUC-01)

[20] J. Morris, G. Hirst,. Lexical cohesion, the thesaurus, and the structure of text. *Computational linguistics*, Vol.17, Num. 1, March 1991, pags. 21-48

[21] S. Muresan, E. Tzoukermann, J. Klavans. 2001 Combining Linguistic and Machine Learning Techniques for Email Summarization. In *ACL-EACL'01 CoNLL Workshop*

[22] M. Palomar, A. Ferrández, L. Moreno, P. Martínez-Barco, J. Peral, M. Saiz-Noeda, R. Muñoz (2001) An Algorithm for Anaphora Resolution in Spanish Texts, *Computational Linguistics*, Vol. 27, Num. 4,

[23] D.R. Radev, H. Jing, M. Budzikowska. 2000. Centroid-based summarization of multiple documents: clustering, sentence extraction, and evaluation In *ANLP/NAACL Workshop on Summarization*, Seattle, WA.

[24] J. D. Schlesinger, D. J. Baker, R. L. Donaway. 2001. Using Document Features and Statistical Modeling to Improve Query-Based Summarization In *Workshop on Text Summarization in conjunction with the ACM SIGIR Conference 2001*, New Orleans, Louisiana USA,

[25] E. Tzoukermann, S. Muresan, J. Klavans. 2001. GIST-IT: Summarizing Email using Linguistic Knowledge and Machine Learning. In *Proceedings of ACL-EACL'01 HLT/KM Workshop*

[26] M. White, T. Korelsky, C. Cardie, V. Ng, D. Pierce, K. Wagstaff. 2001. Multi-document Summarization via Information Extraction. In *Proceedings of the First International Conference on Human Language Technology Research.*

Acknowledgment

The system presented in this paper has been partially funded by TIC2000-0335-C03-01 (Hermes project)