

Resources for Persuasion

Marco Guerini, Carlo Strapparava, Oliviero Stock

FBK-irst
38050 Povo, Trento, Italy
{guerini, strappa, stock}@fbk.eu

Abstract

This paper presents resources and strategies for *persuasive natural language processing*. After the introduction of a specifically tagged corpus, some techniques for affective language processing and for persuasive lexicon extraction are provided together with prospective scenarios of application.

1. Introduction

Persuasive communication is a concept that is becoming important in NLP. In order to automatically produce and analyze persuasive communication, specific resources and methodologies are needed. The paper is structured as follows: Section 2 gives an overview of key concepts connected to persuasion and briefly describes the state of the art in related areas. Section 3 describes the resource we built for statistical acquisition of persuasive expressions. Section 4 describes how this approach can be used for various persuasive NLP tasks.

2. Persuasion, emotions and related concepts

According to Perelman and Olbrechts-Tyteca (Perelman and Olbrechts-Tyteca, 1969), persuasion is a skill that human beings use - in communication - in order to make their partners perform certain actions or collaborate in various activities. Here below we introduce some related key concepts.

Argumentation and Persuasion. In AI the main approaches focus on the argumentative aspects of persuasion. Still, argumentation is considered as a process that involves “rational elements”, while persuasion includes also elements like emotions. In our view, a better distinction can be drawn considering their different foci of attention: while the former focuses on message correctness (its being a valid argument) the latter is concerned with its effectiveness. The recent area of natural argumentation tries to bridge the two (Reed and Grasso, 2007).

Emotions and Persuasion. Since persuasion includes non-rational elements as well, it is a “superset” of argumentation, but this does not rule out that there is a role for emotion within argumentation (Miceli et al., 2006): through arousal of emotions or through appeal to expected emotions. Indeed, emotional communication has become of increasing interest for Persuasive NL Generation.

Rhetorics. The study of how language can be used effectively. This area of studies concerns the linguistic means of persuasion (one of the main means, but not the only one). This is the area we are focusing on in this paper.

Irony. It refers to the practice of saying one thing whilst meaning another. Irony occurs when a word or phrase has

a surface meaning, but another contradictory meaning beneath the surface. Irony is a widely used rhetorical artifice, especially in advertisement.

Persuasion and NLP. The first attempts in persuasive NLG were based mainly on knowledge-based techniques, e.g. Reiter (Reiter et al., 2003), Zukerman (Zukerman et al., 2000). Nowadays approaches to NLP instead are mainly statistically based, calling for collection of annotated texts. In order to explore persuasive NLG in an unrestricted domain, knowledge-based techniques could not suffice: statistical acquisition from large corpora is needed, especially for tasks such as lexical choice.

Another relevant aspect is the increasing importance of emotion and affect for persuasive NLP, e.g. there is a recent area of research called affective NLG - an annotated bibliography can be found in (Piwek, 2002). Affective NLG widens the scope of Opinion Mining (Carenini et al., 2005; Wilson et al., 2004) that only focuses on polarity (valence) recognition for evaluative language retrieval, regardless of any persuasive use of the text.

3. CORPS: Corpora of tagged Political Speeches

The aim of our research is to adopt persuasive expression mining techniques for persuasive NLG in an unrestricted domain. To our knowledge this is the first attempt to use statistical approaches to persuasive messages generation.

As for emotions, we restrict our focus on valenced expressions (i.e. those that have a positive or negative connotation). So the task of producing affective expressions for persuasive goals is treated as the task of changing appropriately the valence of expressions. Specific resources for persuasive NLG are needed. We collected them according to the message characteristics we want to address:

- long and elaborated texts;
- short, high impact, sentences
- single word messages -e.g. brand names-

At present we collected the following resources:

- A CORpus of tagged Political Speeches (CORPS), as examples of long and elaborated persuasive texts

- A Corpus of labeled advertising or political slogans (*SloGun*), as examples of short, high impact, sentences
- A resource containing sets of similar terms with different valence, called OVVTs (Ordered Vectors of Valenced Terms) for the task of affective variation of persuasive expressions (Guerini et al., 2008).

In this paper we focus on the CORPS resource. In collecting this corpus we relied on the hypothesis that tags about public reaction, such as APPLAUSE, are indicators of hot-spots, where persuasion attempts succeeded or, at least, a persuasive attempt has been recognized by the audience; on this point see (Bull and Noordhuizen, 2000) on mistimed applauses in political speeches. We can then perform specific analyses - and extractions - of persuasive linguistic material that causes the audience reaction. Given the textual nature of the corpus, rhetorical artifices based on prosody and other speech features cannot be addressed.

At present, there are about 900 speeches in the corpus and about 2.2 millions words (see Figure 1 for a survey on main speakers' number of speeches). The speeches are all in native English language, and all represent monological situations (e.g. there is only one speaker addressing an audience). We took this decision since dialogical situations (like in political debates) are not in our current focus of research and pose further problem in labeling and analysis.

These speeches have been collected from the Internet; they are manual transcriptions (already annotated with audience reactions tags) of the recordings. We performed an automatic conversion of these tags, to make them homogeneous in formalism and labeling. For example some discourses contain the tag {BIG-APPLAUSE} while others have {LOUD-APPLAUSE}. See Table 1 for a summary of audience reactions tags and their conversion.

Metadata regarding the speech has also been added (title, event, speaker, date, description; see Table 2 for a complete description of the structure of the speeches).

A special tag COMMENT is used for particular cases, e.g.:

- {COMMENT = "A moment of silence was observed"}
- {COMMENT = "An audience member claps"}
- {COMMENT = "recording interrupted"}

With regard to the problem of inter-annotators and inter-sources agreement it should be noted that:

1. The automatic conversion of audience reactions tags drastically reduces the problem of the heterogeneity in tags vocabularies (in fact various sources were considered in order to collect this corpus). These discrepancies can be virtually eliminated at the analysis stage by further clustering tags into coherent groups of audience reactions (see following sections).

2. Since tags represent audience reactions, in principle there is an "evident" high inter-annotators agreement. In some sense it is the audience itself that "annotates" the corpus.

As for the problem of label informativeness, especially if focusing on the problem of mistimed applauses, it should be noted that there are no explicit annotations on applause duration, delay or similar in this corpus, see for example (Atkinson, 1984); so it is difficult to state if and when there has been a mismatching. Still, we believe that for our purposes this is not a problem, because persuasive dynamics are still presents (an "interruptive applause" indicates that there has been an impact on the audience even if not intended by the speaker, a "delayed applause" indicates that there has been a persuasive attempt which has not been promptly recognized by the audience).

Moreover, given the four categories of mistiming proposed by Bull, at least some cases can be individuated:

- An "isolated applause" is individuated by {COMMENT = "An audience member claps"} (when explicitly recorded by annotators). Obviously this tag is not considered as the tag APPLAUSE.
- An "interruptive applause" is individuated by a fragment of speech where a sentence is broken up by an audience intervention (no End Of Sentence mark, dangling sentence parsing, and usually before the tag there is also a double dash to signal the interruption).
- Also the special cases of speakers interrupting applause can be individuated when the speaker explicitly asks for letting him go on.

Text annotation techniques vary according to the degree of manual intervention involved in the annotation process; the (semi)automatic approach we used to collect CORPS limits the amount of costly, manually annotated data. The procedure involved:

- the use of specific HTML parsing algorithms to extract the meta-data from the web-pages (when large scale and homogeneous corpora were available)
- conversion to make homogeneous the tag names (as mentioned before)
- a manual check for consistency of the final output, e.g.
 - (a) the web sources were not uniformly formatted and
 - (b) annotators made typos in tagging.

4. Uses

The CORPS has been used both for analysis and generation. First, to reduce data sparseness, we use a lemmatiser and a part of speech tagger on the whole corpus, that give for each token in the text the corresponding lemma and pos. So, at the lexical level we considered *lemmata* rather than *tokens*. In our approach we further considered:

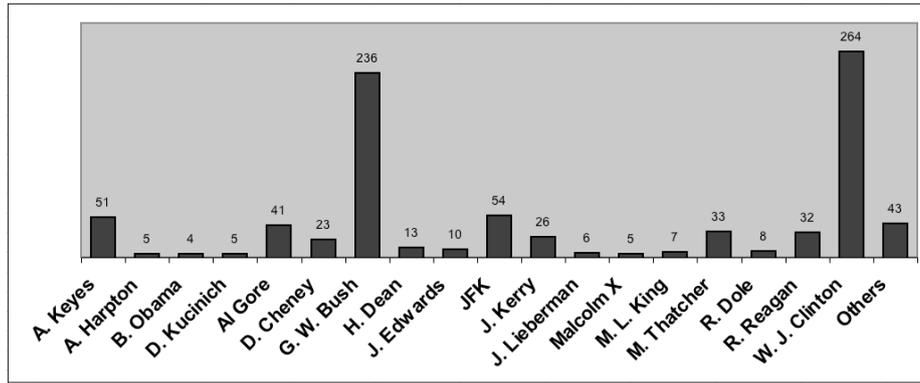


Figure 1: Number of speeches per speaker

Tag	Note
{APPLAUSE}	Main tag in speech transcription.
{SPONTANEOUS-DEMONSTRATION}	Tags replaced: “reaction” “audience interruption”
{STANDING-OVATION}	-
{SUSTAINED-APPLAUSE}	Tags replaced: “big applause” “loud applause” etc.
{CHEERS}	Cries or shouts of approval from the audience. Tags replaced: “cries” “shouts” “whistles” etc.
{BOOING}	In this case, the act of showing displeasure by loudly yelling “Boo” Tags replaced: “hissing”
{TAG1 ; TAG2 ; ...}	In case of multiple tagging, tags are divided by semicolon. Usually there are at most two tags.

Special Tag	Note
{AUDIENCE-MEMBER} [text] {/AUDIENCE-MEMBER}	Tag used to signal a single audience member’s intervention such as clagues speaking.
{OTHER-SPEAK} [text] {/OTHER-SPEAK}	Tag used to signal speakers other than the subject (like journalists, chairmen, etc.)
{AUDIENCE} [text] {/AUDIENCE}	Tag used to signal audience’s intervention.

Table 1: List of main tags

- windows of different width wn (where wn is the number of tokens considered) preceding audience reactions tags
- the typology of persuasive communication (audience

reaction).

As for what concern the last point in the previous list, we individuate three main groups of tags:

{title} [mandatory - describing the speech] {/title}
{event} [not mandatory - derivable from the title] {/event}
{speaker} [mandatory] {/speaker}
{date} [mandatory] {/date}
{source} [mandatory - internet address] {/source}
{description} [if present in the source] {/description}
{speech} [speech transcription with audience reactions tags] {/speech}

Table 2: Structure of a speech entry in CORPS

- *Positive-Focus*: this group indicates a persuasive attempt that sets a positive focus in the audience. Tags considered (about 16 thousand): {APPLAUSE}, {SPONTANEOUS-DEMONSTRATION}, {STANDING-OVATION}, {SUSTAINED-APPLAUSE}, {AUDIENCE-INTERVENTION}, {CHEERING}.
- *Negative-Focus*: It indicates a persuasive attempt that sets a negative focus in the audience. Note that the negative focus is set towards the object of the speech and not on the speaker herself (e.g. “Do we want more taxes?”) Tags considered (about 1 hundred): {BOOING}, {AUDIENCE} NO! {/AUDIENCE}.
- *Ironical*: Indicate the use of ironical devices in persuasion. Tags considered (about 4 thousand): {LAUGHTER}¹.

We conducted a preliminary analysis of the corpus focusing on the relation between valence and persuasion: the phase that leads to audience reaction (e.g. APPLAUSE), if it presents valence dynamics, is characterized by a valence crescendo. That is to say: not necessarily persuasion is achieved via modification of valence intensity, but, when this is the case, it is by means of an increase in the valence of the fragment of speech.

To come to this result we computed, for every window, its mean valence (\bar{w}), calculated by summing up all the valences of the lemmata (SentiWordNet scores) corresponding to the tokens in the fragment and divided by wn , and subtracted the mean valence of the corresponding speech (\bar{s}). In this way we obtained two classes of windows:

- Windows with mean-valence above the mean-valence of the speech ($\bar{w} > \bar{s}$)
- Windows with mean-valence below the mean-valence of the speech ($\bar{s} > \bar{w}$)

We then summed up all the values for the two classes and normalized the results by dividing it for the total number of cases in the class (n_c). We repeated the procedure for

various window widths ($5 < wn < 40$), see Figure 2 and formula 1. The results show that cases above the speech mean are fewer but far stronger. We are planning to have a finer grained analysis by means of cluster-based approaches and variable window width.

$$y = \frac{\sum abs|\bar{w} - \bar{s}|}{n_c} \quad x = wn \quad (1)$$

We then focused on the impact of the lexicon used in the speeches assuming that, for persuasive purposes (both in analysis and generation), not all the words have the same importance. We extracted “persuasive words” by using a coefficient of persuasive impact (pi) based on a weighted tf-idf (see formula 2, $pi = tf \times idf$).

$$tf_i = \frac{n_i \times \sum_{n_i} s_i}{\sum_k n_k} \quad idf_i = \log \frac{|D|}{|\{d : d \ni t_i\}|} \quad (2)$$

To calculate the tf-idf weight, we created a “virtual document” by unifying all the tokens inside all the windows (of dimension $wn = 15$) preceding audience reactions tags, and considering the number of documents in the corpus as coincident to the number of speeches plus one (the virtual document). Obviously from the speeches we subtracted those pieces of text that were used to form the virtual document. Given this premise we can now define the terms in formula 2:

- n_i = number of times the term (word) t_i appears in the virtual document
- $\sum_{n_i} s_i$ = sum of the scores of the word (the closer to the tag the higher the score)
- $\sum_k n_k$ = the number of occurrences of all words = $wn \times |tags\ number|$
- $|D|$ = total number of speeches in the corpus (included the virtual document)
- $|\{d : d \ni t_i\}|$ = number of documents where the term t_i appears (we made an hypothesis of equidistribution).

¹If LAUGHTER appears in a multiple tag (e.g. together with APPLAUSE) by default this tag is associated to the ironical group. This is not the case for BOOING that occurs always alone.

Four lists of words were created according to the group of audience reactions tags they refer to (positive-focus-words,

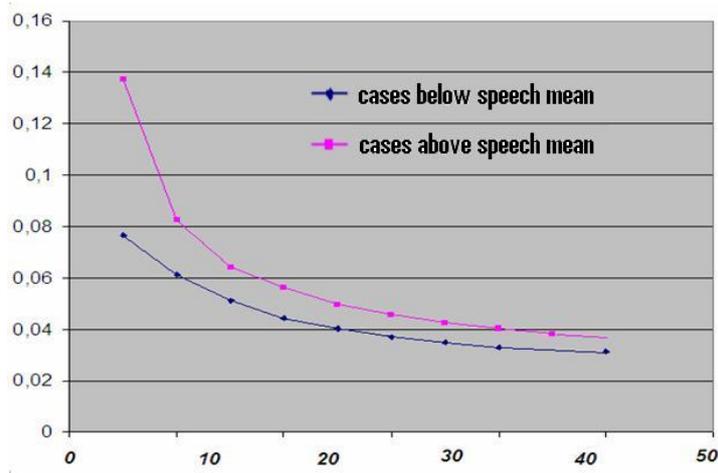


Figure 2: Number of speeches per speaker

negative-focus-words, ironical-words and a persuasive-words list - computed by considering all tags together). Analyzing the 100 top words of these lists (ordered according to their pi score) we found that the negative valence mean of positive-focus and negative-focus groups is the same, while for the negative-focus group the positive valence mean is about 1/4 with regard to the positive-focus group (t-test; $\alpha < 0.01$). These results can be explained by a high use of the CONTRAST relation (that brings negatively valenced words when talking about opponents) in the positive-focus group, while this is not the case for the negative-focus group. In fact a CONTRAST relation is present in about 30% of the sentences preceding APPLAUSE (Atkinson, 1984; Heritage and Greatbatch, 1986).

In Table 3 a comparison between the positive-focus and negative-focus top 50 most persuasive words is given (note that named entities have not been discarded).

It is a matter of debate whether these words are “universally” persuasive (i.e. they could be biased by speaker style, audience typology, context of use and so on). To partially overcome the problem the corpus has been balanced choosing speakers that are equally distributed within the two major parties in the US (Democratic and Republican). At present we do not address the problem of words context (e.g. if a word is preceded by a negation or a hypothetical clause) but we do believe that this does not invalidate the pi of a word. Let us hypothesize that the word `bad#a` has a high pi in the positive-focus list, but the word is mainly used in contexts like “not bad”. This does not imply that the word `bad#a` should be discarded from the positive-focus list, rather that it would be useful to have contextual information for the word (like the co-occurrence score with the word `no#r`).

Analysis of public reaction can substantiate intuitions about the speakers’ rhetorical style. Given the formal annotation of the corpus together with the pi measure we presented, this analysis can be made automatically on a large scale, allowing to gain interesting insights. In fact there are rhetorical phenomena that do not come into light with traditional approaches - based on words *usage* (counting of their occurrences). Considering also words *impact* (their persua-

siveness coefficient pi) a much finer analysis is possible, for example:

How do political speeches change after key historical events? There are works such as (Bligh et al., 2004) that investigated Bush’s speeches lexicon before and after *September 11th* with tools for automatic analysis of political discourses (DICTION 5.0) focusing on charisma traits. Using CORPUS, and analyzing some of the speeches of George W. Bush before and after *September 11th* (70 speeches before and 70 after, from 12 months before to 16 months after) at the lexical level we found that: while the positive valence mean remains totally unvaried, the negative increases by 15% (t-test; $\alpha < 0.001$).

We drove a quantitative/qualitative analysis on Bush’s persuasive words before and after *9/11* to understand how his rhetorics changed (making two lists of persuasive words, one for the speeches before *9/11* and another for the speeches after *9/11*). We focused on some paradigmatic words and found some interesting results. The words are presented in Table 4. In the first column there is the `lemma#pos` (word), in the second and third column its position (persuasiveness)² in the lists before and after *9/11*. On the fourth and fifth columns the number of occurrences in the speeches. An “x” indicates that the word is not “persuasive” (i.e. it appears in the corpus but never in proximity of an audience reaction, the persuasiveness ceases around position 2500 in the lists). An “-” indicates the word is not present in the corpus at all.

A simple approach based on words *usage* was followed in (Bligh et al., 2004). Here, instead, we adopted also words *impact* and created a matrix - for every word - that records an increase/decrease of use compared with an increase/decrease of persuasiveness. Some interesting phenomena emerged. Let us consider the words `military#n` or `treat#v`. Both words are used almost the same number of times before and after *9/11* (respectively 23 vs. 29 times and 25 vs. 20 times). So their “informativeness”, based on number of occurrences, is null. But considering the persuasiveness score we see that their impact vary a lot (respec-

²We use the rank in the list, instead of the pi for readability purposes.

Positive-focus words	Negative-focus words
bless#v deserve#v victory#n justice#n fine#a relief#n November#n win#v help#n thanks#n glad#a stop#v better#r congressman#n lady#n regime#n fabulous#a uniform#n mil- itary#a wrong#a soul#n lawsuit#n welcome#v appreciate#v Bush#n behind#r grateful#a 21st#a defend#v responsible#a safe#a terror#n cause#n bridge#n prevail#v choose#v hand#n love#v frivolous#a sir#n honor#n defeat#v end#v fight#n no#r Joe#n ready#a wear#v future#a direction#n foreign#a death#n single#a democratic#a	horrible#a criticize#v waste#n opponent#n timidity#n shut- tle#n erode#v torpor#n Soviets#n invasion#n scout#n viola- tion#n Castro#n troop#n authority#n Guevara#n Kaufman#n Sachs#n Goldman#n ferociously#r solvent#n page#n front#a international#a direction#n monstrosity#n Cambodia#n un- bearable#a drilling#n Soviet#a increase#v intelligence- gathering#a Carolina#n Gerald#n trusted#a drift#n opera- tion#n WTO#n entry#n mcgovern#v coward#n household#n Neill#n

Table 3: list of top most persuasive words

tively from position 197 to 36 and from 54 to 473). Let us also consider the word `tax#n`, if we consider only the number of occurrences we could infer that before 9/11 this topic was much more “felt” (702 occurrences vs. 81) but if we look at persuasiveness we see that before 9/11 the word `tax#n` never got audiences reactions, while after it become very “popular” (position 93). The same, but in an opposite direction, holds for `war#n`: mentioned three times more after 9/11 (80 vs. 254), but never got an applause.

The results were divided in four blocks, according to thematic areas. In the first block there are words that became very “popular” after 9/11. They usually (indirectly) refer to war and usually from a positive point of view. These words were not considered before 9/11 (i.e. `justice#n` was not persuasive at all before 9/11 and jumped to the ninth position after; at the same time its frequency increased by ten times after the attack). The second block represent words that were “popular” before the attack but became “unutterable” after 9/11 (e.g. `death#n` that fell from position 4 to 450 with an halving in frequency). These words generally refer to the negative aspects of war or to war itself. The third block contains some words that well represent the shifting in the political agenda before and after 9/11 (taxation, contrasting drugs use, leadership). The fourth block shows some abstract and moving words, that became less used and popular after 9/11, partially in contrast with the findings of (Bligh et al., 2004).

What can be said of the lexical choices of a specific speaker that obtains a certain characteristic pattern of public reaction? Ronald Reagan’s (also know as “the great communicator”) rhetorics has been in the focus of many qualitative researches: e.g. interview to ghostwriters, (Collier, 2006), also focusing on particular aspects of his stile like irony, e.g. (Weintraub, 1986) and (Stevenson, 2004). We tried to test whether these findings where consistent with our corpus. By considering 32 of Ronald Reagan’s speeches we first found that the mean tag density of this collection is 1/2 of the mean tag density of the whole corpus (t-test; $\alpha < 0.001$). At first sight this result is somewhat strange, because his being a “great communicator” is not bound to his “firing up” rate (far below the average rate of others speakers). But interestingly, focusing only on the subgroup of ironical tags we found that the density in Reagan’s speeches is almost double as compared to the whole corpus (t-test; $\alpha < 0.001$). The results are even more striking if comparing the mean ironical-tags ratio mtr_i (the mean of the ratio of ironical tags to positive-focus and negative-focus tags per speech, see Formula 3) of the two groups. In Reagan’s

speeches the mtr_i is about 7.5 times greater than the mtr_i of the whole corpus (about 3.5 vs. about 0.5; t-test; $\alpha < 0.001$). That is to say, while normally there is one tags of LAUGHTER every two other tags such as APPLAUSE, in Reagan’s speeches there is one tag such as APPLAUSE out of three, four tags of LAUGHTER.

$$mtr_i = \sum \frac{|ironical-tags|}{|positive-focus| + |negative-focus|} \quad (3)$$

With regard to Reagan’s overall style, his criterion was “Would you talk that way to your barber?”, as reported in (Collier, 2006). He wanted his style to appear “simple and conversational”. To verify this statement we made a hypothesis that a simple and conversational stile is more polysemic than a “cultured” style (richer in technical, and less polysemic, terms). We first calculated the mean polysemy of Reagan’s speeches and compared it to the mean polysemy of the whole corpus, finding no statistical difference between the two (also in this case words use analysis is not informative). Then we focused on the persuasive lexicon: we made a list of Reagan’s persuasive words and compared it to the persuasive words list of the rest of the corpus (we considered all the words whose p_i was different from 0). We found that the mean polysemy of Reagan persuasive words is almost double as compared to the whole corpus (t-test; $\alpha < 0.001$).

How does the perception of the enemy change in different historical moments? A specific analysis on the valence of the lexical context surrounding named entities that elicit negative-focus audience reactions in different period of times can provide interesting insights. Looking at Table 3 it is clear that there are various named entities in the list of negative-focus words at the topmost positions (while this is not the case for positive-focus words). Given the small amount of negative-focus tags our approach will include a second, inductive, analysis step: after individuating named entities that elicit negative-focus reactions (i.e. the “enemies”), those same entities will be searched in the corpus (in the surroundings of positive-focus tags) by assuming that they are inserted in a CONTRAST relation, that sets a temporary negative focus (on enemies behavior), as described before in this section.

Persuasive Opinion Mining. Not all the opinions expressed in speeches or texts have the same persuasive impact. “Successful” opinions (for example G. W. Bush speaking about W. J. Clinton) can be extracted considering those followed by a reaction of the audience. The role of rhetorical constructs will be taken into account in future research.

Lemma	Ranking before	Ranking after	Occurrences before	Occurrences after
win#v	112	7	27	52
justice#n	x	9	15	111
prevail#v	x	15	2	20
defeat#v	x	16	1	44
right#r	x	25	94	55
taliban#n	x	27	1	44
mighty#a	615	30	4	26
military#n	197	36	23	29
victory#n	826	65	9	26
wherever#r	x	115	4	45
evil#a	-	129	0	44
death#n	4	450	65	32
war#n	36	x	80	258
treat#v	54	473	25	20
soldier#n	70	296	20	47
tax#n	x	93	702	81
refund#n	15	-	10	0
wage#n	121	-	4	0
drug-free#a	87	x	9	3
commander-in-chief#n	76	850	25	14
leadership#n	81	261	40	75
future#n	83	394	54	51
dream#n	99	321	77	30
soul#n	23	126	47	32
generation#n	122	442	27	56

Table 4: Bush’s words before and after *September 11th*. In the second and third column, the number represents the rank in the list of persuasive words; an “x” indicates a $pi = 0$; an “-” indicates the word is not present in the corpus at all. In the fourth and fifth columns the total number of occurrences.

5. Conclusions and future work

We have presented the corpus CORPS that contains political speeches tagged with audience reactions. CORPS is freely available for research purposes (for further details see <http://hlt.fbk.eu/corps>) and we want to promote its scaling up. Along with the corpus we have described techniques for statistical acquisition of persuasive expressions (such as a measure of persuasive impact of words) with a view towards contributing to various persuasive NLP tasks.

Regarding lexical choice in text generation micro-planning, there are approaches, e.g. (Jing, 1998), which use corpus and domain information for choosing appropriate lemmata inside synsets. For persuasive NLG, the lists of words we collected allow us to decide, given a synset and an affective/persuasive goal, which lemma to choose, to maximize the impact of the message. We implemented this approach in *Valentino* (Guerini et al., 2008).

Valentino is a tool for increasing the persuasive impact of existing expressions *via* valence modification of the original expression. The novel expressions can then replace previous expressions in the text. *Valentino* uses a term extraction and transformation approach: given a term in the text to be modified, the system accesses the OVVT (briefly mentioned in Section 3) containing that term and chooses the most appropriate transformation in agreement with the valence shift for the persuasive goal and the pi of the words candidates for substitution.

In a complete application, when applied to a text, changes invoked by the strategic level can be of different types. For instance they can uniformly shift to the negative or to the positive polarity, or they can smooth all emotional peaks, or they can strengthen both positive and negative valence. More in general they can be introduced in combination with deeper rhetorical structure analysis, resulting in different types of changes for key parts of the texts.

6. References

- J.M. Atkinson, 1984. *Structures of Social Action*, chapter Public speaking and audience response: some techniques for inviting applause, pages 370–409. Cambridge: CUP.
- M. C. Bligh, J. C. Kohles, and J. R. Meindl. 2004. Charisma under crisis: Presidential leadership, rhetoric, and media responses before and after the september 11th terrorist attacks. *The Leadership Quarterly*, 15(2):211–239.
- P. Bull and M. Noordhuizen. 2000. The mistiming of applause in political speeches. *Journal of Language and Social Psychology*, 19:275–294.
- G. Carenini, R. Ng, and E. Zwart. 2005. Extracting knowledge from evaluative text. In *Proceedings of the 3rd international conference on Knowledge Capture*, pages 11–18.
- K. Collier. 2006. Writing for the great communicators: Speechwriting for roosevelt and reagan. In *Proceedings*

- of the Southwest Political Science Association Meetings, San Antonio, Texas.
- M. Guerini, O. Stock, and C. Strapparava. 2008. Valentino: A tool for valence shifting of natural language texts. In *Proceedings of LREC 2008*, Marrakech, Morocco.
- J. Heritage and D. Greatbatch. 1986. Generating applause: a study of rhetoric and response at party political conferences. *American Journal of Sociology*, 92:110–157.
- H. Jing. 1998. Usage of wordnet in natural language generation. In S. Harabagiu, editor, *Proceedings of the conference "Use of WordNet in Natural Language Processing Systems"*, pages 128–134, Somerset, New Jersey. Association for Computational Linguistics.
- M. Miceli, F. deRosis, and I. Poggi. 2006. Emotional and non-emotional persuasion. *Applied Artificial Intelligence*, 20.
- C. Perelman and L. Olbrechts-Tyteca. 1969. *The new Rhetoric: a treatise on Argumentation*. Notre Dame Press.
- P. Piwek. 2002. An annotated bibliography of affective natural language generation. ITRI ITRI-02-02, University of Brighton.
- C.A. Reed and F. Grasso. 2007. Recent advances in computational models of argument. *International Journal of Intelligent Systems*, 22(1):1–15.
- E. Reiter, S. Sripada, and R. Robertson. 2003. Acquiring correct knowledge for natural language generation. *Journal of Artificial Intelligence Research*, 18:491–516.
- D. J. Stevenson. 2004. Laughter and leadership. presented at The International Center for Studies in Creativity Buffalo State College.
- W. Weintraub. 1986. Personality profiles of american presidents as revealed in their public statements: The presidential news conferences of jimmy carter and ronald reagan. *Political Psychology*, 7(2):285–295.
- T. Wilson, J. Wiebe, and R. Hwa. 2004. Just how mad are you? finding strong and weak opinion clauses. In *Proceedings of AAAI*, pages 761–769.
- I. Zukerman, R. McConachy, and K. Korb. 2000. Using argumentation strategies in automated argument generation. In *Proceedings of the 1st International Natural Language Generation Conference*, pages 55–62.